

**Installation
and
Configuration
Guide**

HP StorageWorks HSG80 ACS Solution Software V8.8

for IBM AIX

Product Version: 8.8-1

First Edition (March 2005)

Part Number: AA-RV1HA-TE

This guide provides installation and configuration instructions and reference material for operation of the HSG80 ACS Solution Software V8.8-1 for IBM AIX.



© Copyright 2000-2005 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Compaq Computer Corporation is a wholly-owned subsidiary of Hewlett-Packard Company.

Intel® and Itanium® are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license.

Microsoft®, MS-DOS®, MS Windows®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

HSG80 ACS Solution Software V8.8 for IBM AIX Installation and Configuration Guide

First Edition (March 2005)

Part Number: AA-RV1HA-TE

contents

About this Guide.....	11
Overview.....	12
Intended Audience	12
Related Documentation	12
Chapter Content Summary	14
Conventions	16
Document Conventions	16
Symbols in Text	16
Symbols on Equipment	17
Rack Stability	19
Getting Help	20
HP Technical Support	20
HP Storage Web Site	20
HP Authorized Reseller	21
Configuration Flowchart.....	22
1 Planning a Subsystem.....	27
Defining Subsystems	28
Controller Designations A and B	28
Controller Designations “This Controller” and “Other Controller”	29
What is Failover Mode?	31
Transparent Failover Mode	31
Multiple-Bus Failover Mode	34
Selecting a Cache Mode	37
Read Caching.....	37
Read-Ahead Caching	37
Write-Back Caching.....	37
Write-Through Caching	38
Enabling Mirrored Caching	39
What is the Command Console LUN?.....	39

Determining the Address of the CCL.....	40
Enabling/Disabling the CCL in SCSI-2 Mode.....	40
Enabling the CCL	40
Disabling the CCL.....	40
Enabling/Disabling CCL in SCSI-3 Mode.....	41
Determining Connections.....	41
Naming Connections	41
Numbers of Connections	42
Assigning Unit Numbers	46
Matching Units to Host Connections in Transparent Failover Mode	46
Matching Units to Host Connections in Multiple-Bus Failover Mode	48
Assigning Unit Numbers Depending on SCSI_VERSION	49
Assigning Host Connection Offsets and Unit Numbers in SCSI-3 Mode.....	49
Assigning Host Connection Offsets and Unit Numbers in SCSI-2 Mode.....	49
What is Selective Storage Presentation?.....	51
Restricting Host Access in Transparent Failover Mode	51
Restricting Host Access by Separate Links.....	52
Restricting Host Access by Disabling Access Paths.....	53
Restricting Host Access by Offsets	55
Restricting Host Access in Multiple-Bus Failover Mode	55
Enabling the Access Path of Selected Host Connections	55
Restricting Host Access by Offsets	57
Worldwide Names (Node IDs and Port IDs).....	58
Restoring Worldwide Names (Node IDs)	59
Unit Worldwide Names (LUN IDs).....	60
2 Planning Storage Configurations	61
Where to Start.....	62
Determining Storage Requirements	64
Configuration Rules for the Controller.....	65
Addressing Conventions for Device PTL.....	67
Examples - Model 2200 Storage Maps, PTL Addressing	69
Choosing a Container Type	76
Creating a Storageset Profile	78
Planning Considerations for Storageset	80
Stripeset Planning Considerations	80
Keep the following points in mind as you plan your stripesets:	81
Mirrorset Planning Considerations	82
Keep these points in mind when planning mirrorsets:	83

RAIDset Planning Considerations	84
Keep these points in mind when planning RAIDsets:.....	85
Striped Mirrorset Planning Considerations	86
Storageset Expansion Considerations	88
Partition Planning Considerations	88
Defining a Partition	89
Guidelines for Partitioning Storagesets and Disk Drives	89
Keep these points in mind when planning partitions for storagesets and disks:.	89
Changing Characteristics Through Switches	90
Enabling Switches	90
Changing Switches.....	90
Specifying Storageset and Partition Switches	91
RAIDset Switches	91
Mirrorset Switches.....	91
Partition Switches.....	92
Specifying Initialization Switches	93
Chunk Size.....	93
Increasing the Request Rate	94
Increasing Sequential Data Transfer Performance	95
Save Configuration.....	95
Destroy/Nodestroy	96
Geometry	96
Specifying Unit Switches	97
Creating Storage Maps	98
Using LOCATE Command to Find Devices	98
Example Storage Map—Model 4310R Disk Enclosure	99
3 Preparing the Host System	103
Installing RAID Array Storage System	104
Making a Physical Connection.....	108
Preparing to Install Host Bus Adapter	108
Installing Host Bus Adapter	108
HP StorageWorks Solution Software Installation	110
New Installation	110
Upgrading Solution Software.....	118
Removing the Solution Software	120
Solution Software Upgrade Procedures	121

4	Installing and Configuring HSG Agent	123
	Why Use StorageWorks Command Console (SWCC)?	124
	Installation and Configuration Overview	126
	About the Network Connection for the Agent	127
	Before Installing the Agent	129
	page 113 Installing the Agent	129
	Event Log File	129
	Configuration Options for the SWCC Agent	130
5	FC Configuration Procedures	133
	Establishing a Local Connection	134
	Setting Up a Single Controller	135
	Powering On and Establishing Communication	135
	Cabling a Single Controller	135
	Configuring a Single Controller Using CLI	136
	Establishing Controller Connection	136
	Configuring Controller Settings	137
	Setting This Controller	137
	Setting Time and Verifying All Commands	140
	Plugging in the FC Cable and Verifying Connections	143
	Repeating Procedure for Each Host Adapter	144
	Setting Up a Controller Pair	145
	Powering Up and Establishing Communication	145
	Cabling a Controller Pair	145
	Configuring a Controller Pair Using CLI	146
	Establishing a Connection to the System Controller	146
	Configuring Controller Settings	147
	Setting This Controller	148
	Setting Time and Verifying All Commands	149
	Plugging in the FC Cable and Verifying Connections	151
	Repeating Procedure for Each Host Adapter Connection	151
	Configuring Devices	152
	Configuring Storage Containers	153
	Configuring a Stripeset	154
	Configuring a Mirrorset	154
	Configuring a RAIDset	155
	Configuring a Striped Mirrorset	156
	Configuring a Single-Disk Unit (JBOD)	157
	Configuring a Partition	157

Assigning Unit Numbers and Unit Qualifiers	159
Assigning a Unit Number to a Storageset	159
Assigning a Unit Number to a Single (JBOD) Disk	159
Assigning a Unit Number to a Partition.....	159
Configuration Options	160
Changing the CLI Prompt	160
Mirroring Cache.....	160
Adding Disk Drives	160
Adding a Disk Drive to the Spareset	160
Removing a Disk Drive from the Spareset	161
Enabling Autospare	161
Deleting a Storageset	162
Changing Switches for a Storageset or Device	162
Displaying the Current Switches.....	163
Changing RAIDset and Mirrorset Switches	163
Changing Device Switches	163
Changing Initialize Switches.....	163
Changing Unit Switches	163
Verifying Storage Configuration from Host.....	164
6 Using CLI for Configuration	165
CLI Configuration Example.....	168
7 Backing Up, Cloning, and Moving Data	177
Backing Up Subsystem Configurations	178
Creating Clones for Backup	179
Clone and Snapshot and AIX Environment	183
Snapshot Tools from the AIX Solution Kit	184
List Snapshot Units	184
Automatic Snapshot.....	184
Manual Snapshot	184
Delete Snapshot.....	185
Recreate Volume Group of a Snapshot	185
Show Disks with the Same Physical Volume ID	185
Recreate Volume Group	185
Manual Snapshot/Clone Creation	186
Moving Storagesets	191

A Subsystem Profile Templates	195
StorageSet Profile	196
Storage Map Template 1 for the BA370 Enclosure	198
Storage Map Template 2 for the Second BA370 Enclosure	199
Storage Map Template 3 for the Third BA370 Enclosure	200
Storage Map Template 4 for the Model 4214R Disk Enclosure	201
Storage Map Template 5 for the Model 4254 Disk Enclosure	203
Storage Map Template 6 for the Model 4310R Disk Enclosure	205
Storage Map Template 7 for the Model 4350R Disk Enclosure	208
Storage Map Template 8 for the Model 4314R Disk Enclosure	210
Storage Map Template 9 for the Model 4354R Disk Enclosure	213
B Installing, Configuring, and Removing the Client	215
Why Install the Client?	216
Before You Install the Client	217
Installing the Client	218
Installing the Integration Patch	219
Should I Install the Integration Patch?	219
Installing the Integration Patch	219
Integrating Controller's SWCC Storage Window with CIM	220
“Insight Manager Unable to Find Controller’s Storage Window”	220
Troubleshooting Client Installation	222
Invalid Network Port Assignments During Installation	222
“There is no disk in the drive” Message	223
Adding Storage Subsystem and its Host to Navigation Tree	224
Removing Command Console Client	226
Where to Find Additional Information	228
About the User Guide	228
About the Online Help	228
C HACMP Notes for AIX	229
Supported Versions of HACMP	229
Installing Solution Software for AIX V8.8 on HACMP Cluster	229
Uninstalling Solution Software for AIX on HACMP Cluster	229
Adding a Virtual Disk After IPL to HACMP Cluster	229
Creating a Volume Group on a Virtual Disk	229

Glossary **235****Index** **257****Figures**

1	General configuration flowchart (panel 1)	23
2	General configuration flowchart (panel 2)	24
3	Configuring storage with SWCC flowchart (panel 3)	25
4	Location of controllers and cache modules in a Model 2200 enclosure	28
5	Location of controllers and cache modules in a BA370 enclosure	29
6	“This controller” and “other controller” for the Model 2200 enclosure	30
7	“This controller” and “other controller” for the BA370 enclosure	30
8	Transparent failover—normal operation	33
9	Transparent failover—after failover from Controller B to Controller A	34
10	Typical multiple-bus configuration	36
11	Mirrored caching	39
12	Connections in separate-link, transparent failover mode configurations	43
13	Connections in single-link, transparent failover mode configurations	44
14	Connections in multiple-bus failover mode	45
15	LUN presentation to hosts, as determined by offset	47
16	Limiting host access in transparent failover mode	53
17	Limiting host access in multiple-bus failover mode	56
18	Placement of the worldwide name label on the Model 2200 enclosure	59
19	Placement of the worldwide name label on the BA370 enclosure	60
20	Mapping a unit to physical disk drives	67
21	PTL naming convention	68
22	How data is laid out on disks in BA370 enclosure configuration	68
23	Storage container types	76
24	Three-member RAID 0 stripeset (example 1)	80
25	Three-member RAID 0 stripeset (example 2)	81
26	Mirrorsets maintain two copies of the same data	83
27	Mirrorset example 2	83
28	Five-member RAIDset using parity	85
29	Striped mirrorset (example 1)	87
30	Striped mirrorset (example 2)	87
31	One example of a partitioned single-disk unit	88
32	Large chunk size increases request rate	94
33	Dual-bus enterprise storage RAID array storage system	106

34	Single-bus enterprise storage RAID array storage system	107
35	An example of a network connection	128
36	Maintenance port connection	134
37	Single controller cabling	136
38	Controller pair failover cabling	146
39	Storage container types	153
40	Example storage map for the BA370 enclosure	166
41	Example, three non-clustered host systems	167
42	Example, logical or virtual disks comprised of storagesets	168
43	Example, virtual system layout from hosts' point of view	168
44	Clone utility steps for duplicating unit members	179
45	Navigation window	224
46	Navigation window showing storage host system "Atlanta"	225
47	Navigation window showing expanded "Atlanta" host icon	225

Tables

1	Summary of chapter contents	14
2	Document conventions	16
3	Unit assignments and SCSI_VERSION	50
4	PTL addressing, single-bus configuration, six Model 4310R enclosures	70
5	PTL addressing, dual-bus configuration, three Model 4350R enclosures	72
6	PTL addressing, single-bus configuration, six Model 4314R enclosures	73
7	PTL addressing, dual-bus configuration, three Model 4354A enclosures	75
8	Comparison of container types	77
9	Example of storageset profile	78
10	Example chunk sizes	95
11	Model 4310R disk enclosure, example of storage map	100
12	SWCC features and components	124
13	Installation and configuration overview	126
14	Correct controller parameter settings	137
15	Correct controller parameter settings, AIX	147

about this guide

This installation guide for HSG80 ACS Solution Software V8.8-1 for IBM AIX provides information to help you:

- Plan the storage array subsystem.
- Install and configure the storage array subsystem on individual operating system platforms.

“About this Guide” topics include:

- ["Overview"](#), page 12
- ["Chapter Content Summary"](#), page 14
- ["Conventions"](#), page 16
- ["Rack Stability"](#), page 19
- ["Getting Help"](#), page 20
- ["HP Authorized Reseller"](#), page 21
- ["Configuration Flowchart"](#), page 22

This book does not contain information about the operating environments to which the controller may be connected, nor does it contain detailed information about subsystem enclosures or their components. See the documentation that accompanied these peripherals for information about them.

Overview

This section covers the following topics:

- ["Intended Audience"](#), page 12
- ["Related Documentation"](#), page 12

Intended Audience

This book is intended for use by systems administrators and systems technicians who are experienced with the following:

- Storage
- Networking

Related Documentation

In addition to this guide, HP provides corresponding information:

- ACS V8.8-1 controller documentation
- *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide*, EK-G80CL-RA. C01
- *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Maintenance and Service Guide*, EK-G80MS-SA. C01
- *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Troubleshooting Guide*, EK-G80TS-SA. C01
- SWCC v2.5 documentation (Client software delivered in solutions kits)
- *HP StorageWorks Command Console User Guide*, AA-RV1UA-TE
- *HP StorageWorks Command Console Release Notes*, AV-RV1TA-TE
- *HP StorageWorks Command Console Help Files*, AA-RS20A-TE and AA-RS21A-TE
- Host-specific documentation (SWCC Agent and HBA software delivered in solutions kits)
- Solution Software Release Notes (platform-specific)

Solution software host support includes the following platforms:

- IBM AIX
- HP-UX
- Linux (Red Hat x86, SuSE x86)
- Novell NetWare
- Open VMS
- Sun Solaris
- Tru64 UNIX
- Windows NT/2000/Windows Server 2003 (32-bit)

Additional support required by HSG80 ACS Solution Software V8.8-1, but delivered through external programs, includes the following:

- Heterogeneous “rules based” SAN configurations
- Host-Bus Adapter (HBA) products
- Applicable Storage Utility Management Suite (SUMS) components
- Vendor-specific switch products
- Secure Path products (Windows, NetWare, Sun, AIX, HP-UX)
- Data Replication Manager (DRM) under ACS 8.8-1
- Enterprise Volume Manager (EVM) under ACS 8.8-1
- Enterprise Backup Solution (EBS)
- Additional ACS variants (W, R)

Chapter Content Summary

Table 1 below summarizes the content of the chapters.

Table 1: Summary of chapter contents

Chapters	Descriptions
1. Planning a Subsystem	This chapter focuses on technical terms and knowledge needed to plan and implement storage array subsystems.
2. Planning Storage Configurations	Plan the storage configuration of your subsystem, using individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives. This chapter describes addressing conventions, configuration rules, creating storage profiles, and creating storage maps.
3. Preparing the Host System	How to prepare your IBM AIX host computer to accommodate the HSG80 controller storage subsystem.
4. Installing and Configuring HSG Agent	The Agent for HSG for a specific operating system polls the storage.
5. FC Configuration Procedures	How to configure a subsystem that uses Fibre Channel (FC) fabric topology. In fabric topology, the controller connects to its hosts through hubs. After the HP StorageWorks Solution Software has been installed, logical units may be created on the RAID Array and configured into the system. Be sure the Agent and Client are installed before logical units are created.
6. Using CLI for Configuration	How-to example of configuring a storage subsystem using the Command Line Interpreter (CLI).

Table 1: Summary of chapter contents (Continued)

Chapters	Descriptions
7. Backing Up, Cloning, and Moving Data	<p>Description of common procedures that are not mentioned elsewhere in this guide.</p> <ul style="list-style-type: none"> ■ Backing Up Subsystem Configuration ■ Cloning Data for Backup ■ Creating Clones and Snapshots for Volume Groups ■ Moving StorageSets
Appendix A. Subsystem Profile Templates	<p>This appendix contains storageSet profiles to copy and use to create your system profiles. It also contains an enclosure template to use to help keep track of the location of devices and storageSets in your shelves. Four (4) templates are needed for the subsystem.</p>
Appendix B. Installing, Configuring, and Removing the Client	<p>The Client monitors and manages a storage subsystem.</p> <p>The following information is included in this appendix:</p> <ul style="list-style-type: none"> ■ Why Install the Client? ■ Before You Install the Client ■ Installing the Client ■ Installing the Integration Patch ■ Troubleshooting Client Installation ■ Adding Storage Subsystem and its Host to the Navigation Tree ■ Removing the Command Console Client ■ Where to Find Additional Information
Appendix C, HACMP Notes for AIX	<p>High availability cluster multi-processing provides install, uninstall and information on adding a virtual disk to HACMP Cluster.</p>

Conventions

Conventions consist of the following:

- "Document Conventions"
- "Symbols in Text"
- "Symbols on Equipment"

Document Conventions

This document follows the conventions in [Table 2](#).

Table 2: Document conventions

Convention	Element
Blue text: Figure 1	Cross-reference links
Bold	Menu items, buttons, and key, tab, and box names
<i>Italics</i>	Text emphasis and document titles in body text
Monospace font	User input, commands, code, file and directory names, and system responses (output and messages)
<i>Monospace, italic font</i>	Command-line and code variables
Blue underlined sans serif font text (http://www.hp.com)	Web site addresses

Symbols in Text

The following symbols may be found in the text of this guide. They have the following meanings:



WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.



Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Tip: Text in a tip provides additional help to readers by providing nonessential or optional techniques, procedures, or shortcuts.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Symbols on Equipment

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings:



Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal injury from electrical shock hazards, do not open this enclosure.



Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal injury from a hot component, allow the surface to cool before touching.



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

Rack stability protects personnel and equipment.



WARNING: To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
 - The full weight of the rack rests on the leveling jacks.
 - In single rack installations, the stabilizing feet are attached to the rack.
 - In multiple rack installations, the racks are coupled.
 - Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.
-

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our web site <http://www.hp.com>.

HP Technical Support

Telephone numbers for worldwide technical support are listed on the following HP web site <http://www.hp.com/support/>. From this web site, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Storage Web Site

The HP web site has the latest information on this product, as well as the latest drivers. Access the storage web site at <http://www.hp.com/country/us/eng/prodserv/storage.html>. From this web site, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP web site for locations and telephone numbers
<http://www.hp.com>.

Configuration Flowchart

A three-part flowchart (Figures 1-3) is shown on the following pages. Refer to these charts while installing and configuring a new storage subsystem. All references in the flowcharts pertain to pages in this guide, unless otherwise indicated.




Figure 1: General configuration flowchart (panel 1)




Figure 2: General configuration flowchart (panel 2)




Figure 3: Configuring storage with SWCC flowchart (panel 3)

Planning a Subsystem

This chapter provides information that helps you plan how to configure the storage array subsystem. This chapter focuses on the technical terms and knowledge needed to plan and implement storage subsystems.

Note: This chapter frequently references the command line interface (CLI). For the complete syntax and descriptions of the CLI commands, see the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide*.

The following information is included in this chapter:

- ["Defining Subsystems"](#), page 28
- ["What is Failover Mode?"](#), page 31
- ["Selecting a Cache Mode"](#), page 37
- ["Enabling Mirrored Caching"](#), page 39
- ["What is the Command Console LUN?"](#), page 39
- ["Determining Connections"](#), page 41
- ["Assigning Unit Numbers"](#), page 46
- ["What is Selective Storage Presentation?"](#), page 51

Note: DILX should be run for ten minutes on all units to delete the 8 MB EISA partition. Refer to *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide* for details.

Refer to [Chapter 2](#) when planning the types of storage containers you need. Storage containers are individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives.

Defining Subsystems

This section describes the terms *this controller* and *other controller*. It also presents graphics of the Model 2200 and BA370 enclosures.

Note: The HSG80 controller uses the BA370 or Model 2200 enclosure.

Controller Designations A and B

The terms A, B, “this controller,” and “other controller,” respectively are used to distinguish one controller from another in a two-controller (also called dual-redundant) subsystem.

Controllers and cache modules are designated either A or B depending on their location in the enclosure, as shown in [Figure 4](#) for the Model 2200 enclosure and in [Figure 5](#) for the BA370 enclosure.




Figure 4: Location of controllers and cache modules in a Model 2200 enclosure





Figure 5: Location of controllers and cache modules in a BA370 enclosure

Controller Designations “This Controller” and “Other Controller”

Some CLI commands use the terms “this” and “other” to identify one controller or the other in a dual-redundant pair. These designations are a shortened form of “this controller” and “other controller.” These terms are defined as follows:

- “**This controller**”—the controller that is the focus of the CLI session. “This controller” is the controller to which the maintenance terminal is attached and through which the CLI commands are being entered. “This controller” can be abbreviated to “this” in CLI commands.
- “**Other controller**”—the controller that is not the focus of the CLI session and through which CLI commands are not being entered. “Other controller” can be abbreviated to “other” in CLI commands.

[Figure 6](#) shows the relationship between “this controller” and “other controller” in a Model 2200 enclosure. [Figure 7](#) shows the same relationship in a BA370 enclosure.




CXO7366A

① This controller

② Other controller

Figure 6: “This controller” and “other controller” for the Model 2200 enclosure



CXO6468D

① Other controller

② This controller

Figure 7: “This controller” and “other controller” for the BA370 enclosure

What is Failover Mode?

Failover is a way to keep the storage array available to the host if one of the controllers becomes unresponsive. A controller can become unresponsive because of a controller hardware failure or, in multiple-bus mode only, due to a failure of the link between host and controller or host-bus adapter. Failover keeps the storage array available to the hosts by allowing the surviving controller to take over total control of the subsystem.

There are two failover modes:

- **Transparent**, which is handled by the surviving controller and is invisible (transparent) to the hosts.
- **Multiple-bus**, which is handled by the surviving controller and which is handled by the hosts through either additional software or as a feature of the operating system.

Either mode of failover can work with loop or fabric topology.

Note: Multiple-bus failover mode is only supported when Secure Path is used. Refer to the Secure Path documentation for more information on using multiple-bus mode.

Transparent Failover Mode

Transparent failover mode has the following characteristics:

- Hosts do not know failover has taken place.
- Units are divided between host ports 1 and 2.

A unit or storageset is a physical or virtual device of the subsystem. It is typically assigned a logical unit number (LUN) and is managed by the HSG80 controller and presented to a server through the Fibre Channel bus and the server's host bus adapter. Disks that are set up as independent disks (JBODs) or RAIDsets are referred to as storagesets.

In transparent failover mode, host port 1 of controller A and host port 1 of controller B must be on the same Fibre Channel link. Host port 2 of controller A and host port 2 of controller B must also be on the same Fibre Channel link. Depending on operating system restrictions and requirements, the port 1 link and the port 2 link can be separate links, or they can be the same link.

At any time, host port 1 is active on only one controller, and host port 2 is active on only one controller. The other ports are in standby mode. In normal operation, both host port 1 on controller A and host port 2 on controller B are active. A representative configuration is shown in [Figure 8](#). The active and standby ports share port identity, enabling the standby port to take over for the active one. If one controller fails, its companion controller (known as the surviving controller) takes control by making both its host ports active, as shown in [Figure 9](#).

Units are divided between the host ports:

- Units 0–99 are on host port 1 of both controllers (but accessible only through the active port).
- Units 100–199 are on host port 2 of both controllers (but accessible only through the active port).

Transparent failover only compensates for a controller failure, and not for failures of either the Fibre Channel link or host Fibre Channel adapters.





Figure 8: Transparent failover—normal operation

CXO7036A



CXO7035A

Figure 9: Transparent failover—after failover from Controller B to Controller A

Multiple-Bus Failover Mode

Multiple-bus failover mode has the following characteristics:

- Host controls the failover process by moving the units from one controller to another
- All units (0 through 199) are visible at all host ports
- Each host has two or more paths to the units

All hosts must have operating system software that supports multiple-bus failover mode. With this software, the host sees the same units visible through two (or more) paths. When one path fails, the host can issue commands to move the units from one path to another. A typical multiple-bus failover configuration is shown in [Figure 10](#).


In multiple-bus failover mode, you can specify which units are normally serviced by a specific controller of a controller pair. Units can be preferred to one controller or the other by the *PREFERRED_PATH* switch of the ADD UNIT (or SET *unit*) command. For example, use the following command to prefer unit D101 to “this controller”:

```
SET D101 PREFERRED_PATH=THIS_CONTROLLER
```

Note: This is an initial preference, which can be overridden by the hosts.

Keep the following points in mind when configuring controllers for multiple-bus failover:

- Multiple-bus failover can compensate for a failure in any of the following:
 - Controller
 - Switch or hub
 - Fibre Channel link
 - Host Fibre Channel adapter
- A host can redistribute the I/O load between the controllers
- All hosts must have operating system software that supports multiple-bus failover mode



NOTE: FCA = Fibre Channel Adapter

CXO7094B

Figure 10: Typical multiple-bus configuration

Selecting a Cache Mode

The cache module supports read, read-ahead, write-through, and write-back caching techniques. The cache technique is selected separately for each unit. For example, you can enable only read and write-through caching for some units while enabling only write-back caching for other units.

Read Caching

When the controller receives a read request from the host, it reads the data from the disk drives, delivers it to the host, and stores the data in its cache module. Subsequent reads for the same data takes the data from cache rather than accessing the data from the disks. This process is called read caching.

Read caching can improve response time to many of the host's read requests. By default, read caching is enabled for all units.

Read-Ahead Caching

During read-ahead caching, the controller anticipates subsequent read requests and begins to prefetch the next blocks of data from the disks as it sends the requested read data to the host. This is a parallel action. The controller notifies the host of the read completion, and subsequent sequential read requests are satisfied from the cache memory. By default, read-ahead caching is enabled for all units.

Write-Back Caching

Write-back caching improves the subsystem's response time to write requests by allowing the controller to declare the write operation complete as soon as the data reaches cache memory. The controller performs the slower operation of writing the data to the disk drives at a later time.

By default, write-back caching is enabled for all units, but only if there is a backup power source for the cache modules (either batteries or an uninterruptible power supply).

Write-Through Caching

Write-through caching is enabled when write-back caching is disabled. When the controller receives a write request from the host, it places the data in its cache module, writes the data to the disk drives, then notifies the host when the write operation is complete. This process is called write-through caching because the data actually passes through—and is stored in—the cache memory on its way to the disk drives.

Enabling Mirrored Caching

In mirrored caching, half of each controller's cache mirrors the companion controller's cache, as shown in [Figure 11](#).

The total memory available for cached data is reduced by half, but the level of protection is greater.




Figure 11: Mirrored caching

Before enabling mirrored caching, make sure the following conditions are met:

- Both controllers support the same size cache.
- Diagnostics indicate that both caches are good.
- No unit errors are outstanding, for example, lost data or data that cannot be written to devices.
- Both controllers are started and configured in failover mode.

What is the Command Console LUN?

StorageWorks Command Console (SWCC) software communicates with the HSG80 controllers through an existing storage unit, or logical unit number (LUN). The dedicated LUN that SWCC uses is called the Command Console LUN (CCL). The CCL serves as the communication device for the HS-Series Agent and identifies itself to the host by a unique identification string. By default, a CCL device is enabled within the HSG80 controller on host port 1. The HSG80 uses SCSI-3 with your platform.

The CCL does the following:

- Allows the RAID Array to be recognized by the host as soon as it is attached to the SCSI bus and configured into the operating system.
- Serves as a communications device for the HS-Series Agent. The CCL identifies itself to the host by a unique identification string.

In dual-redundant controller configurations, the commands described in the following sections alter the setting of the CCL on both controllers. The CCL is enabled only on host port 1. At least one storage device of any type must be configured on host port 2 before installing the Agent on a host connected to host port 2.

Select a storageset that you plan to configure and that is not likely to change. This storageset can be used by the Agent to communicate with the RAID Array. Deleting this storageset (LUN) later breaks the connection between the Agent and the RAID Array.

Determining the Address of the CCL

CCL is enabled by default. Its address can be determined by entering the following CLI command:

```
HSG80 > SHOW THIS_CONTROLLER
```

Enabling/Disabling the CCL in SCSI-2 Mode

Enabling the CCL

To enable the CCL, enter the following CLI command:

```
HSG80 > SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

Disabling the CCL

To disable the CCL, enter the following CLI command:

```
HSG80 > SET THIS_CONTROLLER NOCOMMAND_CONSOLE_LUN
```

To see the state of the CCL, use the SHOW THIS_CONTROLLER/ OTHER_CONTROLLER command. Because the CCL is not an actual LUN, the SHOW UNITS command does not display the CCL location.

Enabling/Disabling CCL in SCSI-3 Mode

The CCL is enabled all the time. There is no option to enable/disable.


Determining Connections

The term “connection” applies to every path between a Fibre Channel adapter in a host computer and an active host port on a controller.

Note: In ACS V8.8-1, the maximum number of supported connections is 96.

Naming Connections

It is highly recommended that you assign names to connections that have meaning in the context of your particular configuration. One system that works well is to name each connection after its host, its adapter, its controller, and its controller host port, as follows:



Examples:

A connection from the first adapter in the host named RED that goes to port 1 of controller A would be called RED1A1.


A connection from the third adapter in host GREEN that goes to port 2 of controller B would be called GREEN3B2.

Note: Connection names can have a maximum of 9 characters.

Numbers of Connections

The number of connections resulting from cabling one adapter into a switch or hub depends on failover mode and how many links the configuration has:


- If a controller pair is in transparent failover mode and the port 1 link is separate from the port 2 link (that is, ports 1 of both controllers are on one loop or fabric, and port 2 of both controllers are on another), each adapter has one connection, as shown in [Figure 12](#).
- If a controller pair is in transparent failover mode and port 1 and port 2 are on the same link (that is, all ports are on the same loop or fabric), each adapter has two connections, as shown in [Figure 13](#).
- If a controller pair is in multiple-bus failover mode, each adapter has two connections, as shown in [Figure 14](#).



NOTE: FCA = Fibre Channel Adapter

CXO7081B


Figure 12: Connections in separate-link, transparent failover mode configurations



NOTE: FCA = Fibre Channel Adapter

CXO7079B

Figure 13: Connections in single-link, transparent failover mode configurations



NOTE: FCA = Fibre Channel Adapter

CXO7080B

Figure 14: Connections in multiple-bus failover mode

Assigning Unit Numbers

The controller keeps track of the unit with the unit number. The unit number can be from 0–199 prefixed by a D, which stands for disk drive. A unit can be presented as different LUNs to different connections. The interaction of a unit and a connection is determined by several factors:

- Failover mode of the controller pair
- The ENABLE_ACCESS_PATH and PREFERRED_PATH switches in the ADD UNIT (or SET *unit*) commands
- The UNIT_OFFSET switch in the ADD CONNECTIONS (or SET *connections*) commands
- The controller port to which the connection is attached
- The SCSI_VERSION switch of the SET THIS_CONTROLLER/OTHER_CONTROLLER command

The considerations for assigning unit numbers are discussed in the following sections.

Matching Units to Host Connections in Transparent Failover Mode

In transparent failover mode, the ADD UNIT command creates a unit for host connection to access and assigns it to either port 1 or port 2 of both controllers.

Unit numbers are assigned to ports as follows:

- 0–99 are assigned to host port 1 of both controllers.
- 100–199 are assigned to host port 2 of both controllers.

For example, unit D2 is on port 1, and unit D102 is available through port 2.

The LUN number that a host connection assigns to a unit is a function of the UNIT_OFFSET switch of the ADD (or SET) CONNECTIONS command. The relationship of offset, LUN number, and unit number is shown in the following equation:

$$\text{LUN number} = \text{unit number} - \text{offset}$$

Where...

- LUN number is relative to the host (what the host sees the unit as)
- Unit number is relative to the controller (what the controller sees the unit as)

If no value is specified for offset, then connections on port 1 have a default offset of 0 and connections on port 2 have a default offset of 100.

For example, if all host connections use the default offset values, unit D2 is presented to a port 1 host connection as LUN 2 (unit number of 2 minus offset of 0). Unit D102 is presented to a port 2 host connection as LUN 2 (unit number of D102 minus offset of 100).

[Figure 15](#) shows how units are presented as different LUNs, depending on the offset of the host. In this illustration, host connection 1 and host connection 2 would need to be on host port 1; host connection 3 would need to be on host port 2.

Controller units	Host connection 1 Offset: 0	Host connection 2 Offset: 20	Host connection 3 Offset: 100
D0	LUN 0		
D1	LUN 1		
D2	LUN 2		
D3	LUN 3		
D20	LUN 20	LUN 0	
D21	LUN 21	LUN 1	
D100			
D101			
D102			
D130			
D131			

CXO6455B

Figure 15: LUN presentation to hosts, as determined by offset

Offsets other than the default values can be specified. For example, unit D17 would be visible to a host connection on port 1 that had an offset of 10 as LUN 7 (unit number of 17 minus offset of 10). The unit would not be visible to a host connection with a unit offset of 18 or greater, because that offset is not within the unit's range (unit number of 17 minus offset of 18 is a negative number).

Similarly, unit D127 would be visible to a host connection on port 2 that had an offset of 120 as LUN 7 (unit number of 127 minus offset of 120). The unit would not be visible to a host connection with a unit offset of 128 or greater, because that offset is not within the unit's range (unit number of 127 minus offset of 128 is a negative number).

An additional factor to consider when assigning unit numbers and offsets is SCSI version. If the SCSL_VERSION switch of the SET THIS_CONTROLLER/OTHER_CONTROLLER command is set to SCSI-3, the CCL is presented as LUN 0 to every connection, superseding any unit assignments. The interaction between SCSI version and unit numbers is explained further in the next section.

In addition, the access path to the host connection must be enabled for the connection to access the unit. See "[Restricting Host Access in Transparent Failover Mode](#)", page 51.

Matching Units to Host Connections in Multiple-Bus Failover Mode

In multiple-bus failover mode, the ADD UNIT command creates a unit for host connections to access. All unit numbers (0 through 199) are potentially visible on all four controller ports, but are accessible only to those host connections for which access path is enabled and which have offsets in the unit's range.

The LUN number a host connection assigns to a unit is a function of the UNIT_OFFSET switch of the ADD (or SET) CONNECTIONS command. The default offset is 0. The relationship of offset, LUN number, and unit number is shown in the following equation:

$$\text{LUN number} = \text{unit number} - \text{offset}$$

Where...

- LUN number is relative to the host (number the host sees the unit as)
- Unit number is relative to the controller (number the controller sees the unit as)

For example, unit D7 would be visible to a host connection with an offset of 0 as LUN 7 (unit number of 7 minus offset of 0). Unit D17 would be visible to a host connection with an offset of 10 as LUN 7 (unit number of 17 minus offset of 10). The unit would not be visible at all to a host connection with a unit offset of 18 or greater, because that offset is not within the units range (unit number of 17 minus offset of 18 is a negative number).

In addition, the access path to the host connection must be enabled for the connection to access the unit. This is done through the ENABLE_ACCESS_PATH switch of the ADD UNIT (or SET *unit*) command.

The PREFERRED_PATH switch of the ADD UNIT (or SET *unit*) command determines which controller of a dual-redundant pair initially accesses the unit. Initially, PREFERRED_PATH determines which controller presents the unit as Ready. The other controller presents the unit as Not Ready. Hosts can issue a SCSI Start Unit command to move the unit from one controller to the other.

Assigning Unit Numbers Depending on SCSI_VERSION

The SCSI_VERSION switch of the SET THIS_CONTROLLER/OTHER_CONTROLLER command determines how the CCL is presented. There are two choices: SCSI-2 and SCSI-3. The choice for SCSI_VERSION affects how certain unit numbers and certain host connection offsets interact.

Assigning Host Connection Offsets and Unit Numbers in SCSI-3 Mode

If SCSI_VERSION is set to SCSI-3, the CCL is presented as LUN 0 to all connections. The CCL supersedes any other unit assignment. Therefore, in SCSI-3 mode, a unit that would normally be presented to a connection as LUN 0 is not visible to that connection at all.

The following methods are recommended for assigning host connection offsets and unit numbers in SCSI-3 mode:

- Offsets should be divisible by 10 (for consistency and simplicity).
- Unit numbers should not be assigned at connection offsets (to avoid being masked by the CCL at LUN 0).

For example, if a host connection has an offset of 20 and SCSI-3 mode is selected, the connection sees LUNs as follows:

LUN 0—CCL
LUN 1—unit 21
LUN 2—unit 22, etc.

In this example, if a unit 20 is defined, it is superseded by the CCL and invisible to the connection.

Assigning Host Connection Offsets and Unit Numbers in SCSI-2 Mode

Some operating systems expect or require a disk unit to be at LUN 0. In this case, it is necessary to specify SCSI-2 mode.

If SCSI_VERSION is set to SCSI-2 mode, the CCL floats, moving to the first available LUN location, depending on the configuration.

HP recommends that you use the following conventions when assigning host connection offsets and unit numbers in SCSI-2 mode:

- Offsets should be divisible by 10 (for consistency and simplicity).
- Unit numbers should be assigned at connection offsets (so that every host connection has a unit presented at LUN 0).

Table 3 summarizes the recommendations for unit assignments based on the SCSI_VERSION switch.

Table 3: Unit assignments and SCSI_VERSION

SCSI_VERSION	Offset	Unit Assignment	What the connection sees LUN 0 as
SCSI-2	Divisible by 10	At offsets	Unit whose number matches offset
SCSI-3	Divisible by 10	Not at offsets	CCL

What is Selective Storage Presentation?

Selective Storage presentation is a feature of the HSG80 controller that enables you to control the allocation of storage space and shared access to storage across multiple hosts. This is also known as *Restricting Host Access*.

In a subsystem that is attached to more than one host or if the hosts have more than one adapter, it is possible to reserve certain units for the exclusive use of certain host connections.

For a controller pair, the method used to restrict host access depends on whether the controllers are in transparent or multiple-bus failover mode. For a single controller, the methods are the same as for a controller pair in transparent failover.

The default condition is `ENABLE_ACCESS_PATH=ALL`. This specifies that access paths to ALL hosts are enabled. HP recommends that you restrict host access and that the access path be carefully specified to avoid providing undesired host connections access to the unit.

Now with V8.8-1 you may change the subsystem behavior so that units are always created without enabling connections. This provides more control in granting appropriate access to specific connections.

The command syntax to disable is:

```
HSG> SET this/other Default_Access=Disable
```

The command syntax to enable is:

```
HSG> SET Default this/other Default_Access=Enable (default after  
upgrade)
```

When the command is invoked from one controller, the `Default_Access` from the other controller will be similarly modified. The setting is symmetrical and persistent across restarts, FRUTIL, etc.

Restricting Host Access in Transparent Failover Mode

Three methods can be used to restrict host access to storage units in transparent failover mode:


- Using separate Fibre Channel links (either loop or fabric)
- Enabling the access path of selected host connections on a shared loop or fabric
- Setting offsets

Note: These techniques also work for a single controller.

Restricting Host Access by Separate Links

In transparent failover mode, host port 1 of controller A and host port 1 of controller B share a common Fibre Channel link. Host port 2 of controller A and host port 2 of controller B also share a common Fibre Channel link. If the host 1 link is separate from the host 2 link, the simplest way to limit host access is to have one host or set of hosts on the port 1 link, and another host or set of hosts on the port 2 link. Each host can then see only units assigned to its respective controller port. This separation of host buses is shown in [Figure 16](#). This method applies only if the host 1 link and host 2 link are separate links.

Note: It is highly recommended that you provide access to only specific connections. This way, if new connections are added, they do not have automatic access to all units. See the following section “[Restricting Host Access by Disabling Access Paths](#)” on page 53.



NOTE: FCA = Fibre Channel Adapter

CXO7081B

Figure 16: Limiting host access in transparent failover mode

Restricting Host Access by Disabling Access Paths

If more than one host is on a link (that is, attached to the same port), host access can be limited by enabling the access of certain host connections and disabling the access of others. This is done through the ENABLE_ACCESS_PATH and DISABLE_ACCESS_PATH switches of the ADD UNIT (or SET *unit*) commands. The access path is a unit switch, meaning it must be specified for each unit. Default access enables the unit to be accessible to all hosts.

For example:

In [Figure 17](#), restricting the access of unit D101 to host 3, the host named BROWN can be done by enabling only the connection to host 3. Enter the following commands:

```
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=BROWN1B2
```

If the storage subsystem has more than one host connection, carefully specify the access path to avoid providing undesired host connections access to the unit. The default condition for a unit is that access paths to all host connections are enabled. To restrict host access to a set of host connections, specify DISABLE_ACCESS_PATH=ALL for the unit, then specify the set of host connections that are to have access to the unit.

Enabling the access path to a particular host connection does not override previously enabled access paths. All access paths previously enabled are still valid; the new host connection is simply added to the list of connections that can access the unit.

Note: The procedure of restricting access by enabling all access paths then disabling selected paths is **not recommended** because of the potential data/security breach that occurs when a new host connection is added.

Restricting Host Access by Offsets

Offsets establish the start of the range of units that a host connection can access.

For example:

In [Figure 16](#), assume both host connections on port 2 (connections BLACK1B2 and BROWN1B2) initially have the default port 2 offset of 100. Setting the offset of connection BROWN1B2 to 120 presents unit D120 to host BROWN as LUN 0.

```
SET BROWN1B2 UNIT_OFFSET=120
```

Host BROWN cannot see units lower than its offset, so it cannot access units D100 and D101. However, host BLACK can still access D120 as LUN 20 if the operating system permits. To restrict access of D120 to only host BROWN, enable only host BROWN's access, as follows:

```
SET D120 DISABLE_ACCESS_PATH=ALL  
SET D120 ENABLE_ACCESS_PATH=BROWN1B2
```

Note: HP recommends that you provide access to only specific connections, even if there is just one connection on the link. This way, if new connections are added, they do not have automatic access to all units.


Restricting Host Access in Multiple-Bus Failover Mode

In multiple-bus mode, the units assigned to any port are visible to all ports. There are two ways to limit host access in multiple-bus failover mode:

- Enabling the access path of selected host connections
- Setting offsets

Enabling the Access Path of Selected Host Connections

Host access can be limited by enabling the access of certain host connections and disabling the access of others. This is done through the ENABLE_ACCESS_PATH and DISABLE_ACCESS_PATH switches of the ADD UNIT (or SET *unit*) commands. Access path is a unit switch, meaning it must be specified for each unit. Default access means that the unit is accessible to all hosts. It is important to remember that at least two paths between the unit and the host must be enabled in order for multiple-bus failover to work.



NOTE: FCA = Fibre Channel Adapter

CX07078

Figure 17: Limiting host access in multiple-bus failover mode

For example:

Figure 17 shows a representative multiple-bus failover configuration. Restricting the access of unit D101 to host BLUE can be done by enabling only the connections to host BLUE. At least two connections must be enabled for

multiple-bus failover to work. For most operating systems, it is desirable to have all connections to the host enabled. To enable all connections for host BLUE, enter the following commands:

```
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=BLUE1A1,BLUE1B1,BLUE2A2,BLUE2B2
```

To enable only two connections for host BLUE (if it is a restriction of the operating system), select two connections that use different adapters, different switches or hubs, and different controllers:

```
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=(BLUE1A1, BLUE2B2)  
or  
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=(BLUE1B1, BLUE2A2)
```

If the storage subsystem has more than one host connection, the access path must be specified carefully to avoid giving undesirable host connections access to the unit. The default condition for a unit is that access paths to all host connections are enabled. To restrict host access to a set of host connections, specify DISABLE_ACCESS_PATH=ALL when the unit is added, then use the SET *unit* command to specify the set of host connections that are to have access to the unit.

Enabling the access path to a particular host connection does not override previously enabled access paths. All access paths previously enabled are still valid; the new host connection is simply added to the list of connections that can access the unit.

Note: The procedure of restricting access by enabling all access paths then disabling selected paths is **not recommended** because of the potential data/security breach that occurs when a new host connection is added.

Restricting Host Access by Offsets

Offsets establish the start of the range of units that a host connection can access. However, depending on the operating system, hosts that have lower offsets may be able to access the units in the specified range.

Note: All host connections to the same host computer must be set to the same offset.

For example:

In [Figure 17](#), assume all host connections initially have the default offset of 0. Giving all connections access to host BLUE, an offset of 120 presents unit D120 to host BLUE as LUN 0. Enter the following commands:

```
SET BLUE1A1 UNIT_OFFSET=120  
SET BLUE1B1 UNIT_OFFSET=120  
SET BLUE2A2 UNIT_OFFSET=120  
SET BLUE2B2 UNIT_OFFSET=120
```

Host BLUE cannot see units lower than its offset, so it cannot access any other units. However, the other two hosts can still access D120 as LUN 20 if their operating system permits. To restrict access of D120 to only host BLUE, enable only host BLUE's access, as follows:

```
SET D120 DISABLE_ACCESS_PATH=ALL  
SET D120 ENABLE_ACCESS_PATH=(BLUE1A1,BLUE1B1,BLUE2A2,BLUE2B2)
```

Note: HP recommends that you always provide access to only specific connections. This way, if new connections are added, they do not have automatic access to all units.

Note: See "[Restricting Host Access by Disabling Access Paths](#)", page 53.

Worldwide Names (Node IDs and Port IDs)

A worldwide name—also called a node ID—is a unique, 64-bit number assigned to a subsystem prior to shipping. The node ID belongs to the subsystem itself and never changes.

Each subsystem's node ID ends in zero, for example 5000-1FE1-FF0C-EE00. The controller port IDs are derived from the node ID.

In a subsystem with two controllers in transparent failover mode, the controller port IDs are incremented as follows:

- Controller A and controller B, port 1—worldwide name + 1, for example 5000-1FE1-FF0C-EE01
- Controller A and controller B, port 2—worldwide name + 2, for example 5000-1FE1-FF0C-EE02

In multiple-bus failover mode, each of the host ports has its own port ID:

- Controller B, port 1—worldwide name + 1, for example
5000-1FE1-FF0C-EE01
- Controller B, port 2—worldwide name + 2, for example
5000-1FE1-FF0C-EE02
- Controller A, port 1—worldwide name + 3, for example
5000-1FE1-FF0C-EE03
- Controller A, port 2—worldwide name + 4, for example
5000-1FE1-FF0C-EE04

Use the CLI command, SHOW THIS_CONTROLLER/OTHER_CONTROLLER to display the subsystem's worldwide name.

Restoring Worldwide Names (Node IDs)

If a situation occurs that requires you to restore the worldwide name, you can restore it using the worldwide name and checksum printed on the sticker on the frame into which the controller is inserted.

[Figure 18](#) shows the placement of the worldwide name label for the Model 2200 enclosure, and [Figure 19](#) shows the placement for the BA370 enclosure.




Figure 18: Placement of the worldwide name label on the Model 2200 enclosure




Figure 19: Placement of the worldwide name label on the BA370 enclosure



Caution: Each subsystem has its own unique worldwide name (node ID). If you attempt to set the subsystem worldwide name to a name other than the one that came with the subsystem, the data on the subsystem is not accessible. Never set two subsystems to the same worldwide name, or data corruption occurs.

Unit Worldwide Names (LUN IDs)

In addition, each unit has its own worldwide name, or LUN ID. This is a unique, 128-bit value that the controller assigns at the time of unit initialization. It cannot be altered by you but does change when the unit is reinitialized. Use the SHOW command to list the LUN ID.

Planning Storage Configurations

2

This chapter provides information to help you plan the storage configuration of your subsystem. Storage containers are individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives. Use the guidelines found in this section to plan the various types of storage containers needed.

This chapter also focuses on the required design and implementation aspects, such as addressing conventions, configuration rules, creating storage profiles, and creating storage maps.

The following storage configuration information can be found in this chapter:

- "Where to Start", page 62
- "Determining Storage Requirements", page 64
- "Configuration Rules for the Controller", page 65
- "Addressing Conventions for Device PTL", page 67
- "Choosing a Container Type", page 76
- "Creating a Storageset Profile", page 78
- "Planning Considerations for Storageset", page 80
- "Changing Characteristics Through Switches", page 90
- "Specifying Storageset and Partition Switches", page 91
- "Specifying Initialization Switches", page 93
- "Specifying Unit Switches", page 97
- "Creating Storage Maps", page 98

Refer to [Chapter 3](#) for instructions on how to prepare your host computer to accommodate the HSG80 controller storage subsystem.

Where to Start

The following procedure outlines the steps to follow when planning your storage configuration.

See [Appendix A](#) to locate the blank templates for keeping track of the containers being configured.

1. Determine your storage requirements. Use the questions in "[Determining Storage Requirements](#)", page 64, to help you.
2. Review configuration rules. See "[Configuration Rules for the Controller](#)", page 65.
3. Familiarize yourself with the current physical layout of the devices and their addressing scheme. See "[Addressing Conventions for Device PTL](#)", page 67.
4. Choose the type of storage containers you need to use in your subsystem. See "[Choosing a Container Type](#)", page 76, for a comparison and description of each type of storageset.
5. Create a storageset profile (described in "[Creating a Storageset Profile](#)", page 78). Fill out the storageset profile while you read the sections that pertain to your chosen storage type:
 - ["Planning Considerations for Storageset"](#), page 80
 - ["Mirrorset Planning Considerations"](#), page 82
 - ["RAIDset Planning Considerations"](#), page 84
 - ["Partition Planning Considerations"](#), page 88
 - ["Striped Mirrorset Planning Considerations"](#), page 86
6. Decide which switches you need for your subsystem. General information on switches is detailed in "[Specifying Storageset and Partition Switches](#)", page 91.
 - Determine the unit switches you want for your units ("[Specifying Unit Switches](#)", page 97).
 - Determine the initialization switches you want for your planned storage containers ("[Specifying Initialization Switches](#)", page 93).
7. Create a storage map ("[Creating Storage Maps](#)", page 98).
8. Configure the storage you have now planned using one of the following methods:
 - Use SWCC. See the SWCC documentation for details.

- Use the Command Line Interpreter (CLI) commands. This method allows you flexibility in defining and naming your storage containers. See the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide*.

Determining Storage Requirements

It is important to determine your storage requirements. Here are a few of the questions you should ask yourself regarding the subsystem usage:

- What applications or user groups access the subsystem? How much capacity do they need?
- What are the I/O requirements? If an application is data transfer-intensive, what is the required transfer rate? If it is I/O request-intensive, what is the required response time? What is the read/write ratio for a typical request?
- Are most I/O requests directed to a small percentage of the disk drives? Do you want to keep it that way or balance the I/O load?
- Do you store mission-critical data? Is availability the highest priority or would standard backup procedures suffice?

Configuration Rules for the Controller

The following list defines maximum configuration rules for the controller:

- 128 visible LUNs/200 assignable unit numbers
 - In SCSI-2 mode, if the CCL is enabled, the result is 127 visible LUNs and one CCL.
 - In SCSI-3 mode, if the CCL is enabled, the result is 126 visible LUNs and two CCLs.

Note: In SCSI-2 mode, if the Command Console LUN (CCL) is enabled, and in SCSI-3 mode where the CCL is always enabled, the result is 127 visible LUNs and one CCL.

- 1.024 TB storageset size
- 96 host connections
- 72 physical devices for BA370-based packaging; 84 physical devices for all others
- 20 RAID 3/5 storagesets
- 30 RAID 3/5 and RAID 1 storagesets (see note)
- 45 RAID 3/5, RAID 1, and RAID 0 storagesets (see note)

Note: For the previous two storageset configurations, this is a combined maximum, limited to no more than 20 RAID 3/5 storagesets in the individual combination.

- 8 partitions of a storageset or individual disk
- 6 physical devices per RAID 1 storageset (mirrorset)
- 14 physical devices per RAID 3/5 storageset (RAIDset)
- 24 physical devices per RAID 0 storageset (stripeset)
- 48 physical devices per RAID 0+1 storageset (striped mirrorset)
- 146 GB disks (for universal drives only) are the largest supported for the HSG60 and HSG80 storage array.

Tip: If you are redeploying disks that have been operating under a prior version of ACS into a newly established container, as a best practice, *always* initialize the devices and the new container before proceeding with subsystem activities to avoid operational and performance issues.

Addressing Conventions for Device PTL

The HSG80 controller has six SCSI device ports, each of which connects to a SCSI bus. In dual-controller subsystems, these device buses are shared between the two controllers. (The StorageWorks Command Console calls the device ports “channels.”) The standard BA370 enclosure provides a maximum of four SCSI target identifications (ID) for each device port. If more target IDs are needed, expansion enclosures can be added to the subsystem. For an example of how units are mapped to physical disk drives, see [Figure 20](#).




Figure 20: Mapping a unit to physical disk drives

The HSG80 controller identifies devices based on a Port-Target-LUN (PTL) numbering scheme, shown in [Figure 21](#). The physical location of a device in its enclosure determines its PTL.

- P—Designates the controller's SCSI device port number (1 through 6).
- T—Designates the target ID number of the device. Valid target ID numbers for a single-controller configuration and dual-redundant controller configuration are 0–3 and 8–15, respectively. (This applies to the BA370 cabinet only.)

- L—Designates the logical unit (LUN) of the device. For disk devices the LUN is always 0.

Disk 10200




Figure 21: PTL naming convention

The controller can either operate with a BA370 enclosure or with a Model 2200 controller enclosure combined with Model 4214R, Model 4254, Model 4310R, Model 4350R, Model 4314R, or Model 4354R disk enclosures.

The controller operates with BA370 enclosures that are assigned ID numbers 0, 2, and 3. These ID numbers are set through the PVA module. Enclosure ID number 1, which assigns devices to targets 4 through 7, is not supported. [Figure 22](#) shows how data is laid out on disks in an extended configuration.




Figure 22: How data is laid out on disks in BA370 enclosure configuration

Examples - Model 2200 Storage Maps, PTL Addressing

The Model 2200 controller enclosure can be combined with the following:

- **Model 4214R disk enclosure**—Ultra2 SCSI with 14 drive bays, single-bus I/O module.
- **Model 4254 disk enclosure**—Ultra2 SCSI with 14 drive bays, dual-bus I/O module.

Note: The Model 4214R uses the same storage maps as the Model 4314R, and the Model 4254 uses the same storage maps as the Model 4354R disk enclosures.

- **Model 4310R disk enclosure**—Ultra3 SCSI with 10 drive bays, single-bus I/O module. [Table 4](#) shows the addresses for each device in a six-shelf, single-bus configuration. A maximum of six Model 4310R disk enclosures can be used with each Model 2200 controller enclosure.

Note: The storage map for the Model 4310R reflects the disk enclosure's physical location in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure, and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.

- **Model 4350R disk enclosure**—Ultra3 SCSI with 10 drive bays, single-bus I/O module. [Table 5](#) shows the addresses for each device in a three-shelf, single-bus configuration. A maximum of three Model 4350R disk enclosures can be used with each Model 2200 controller enclosure.
- **Model 4314R disk enclosure**—Ultra3 SCSI with 14 drive bays, single-bus I/O module. [Table 6](#) shows the addresses for each device in a six-shelf, single-bus configuration. A maximum of six Model 4314R disk enclosures can be used with each Model 2200 controller enclosure.

Note: The storage map for the Model 4314R reflects the disk enclosure's physical location in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure, and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.

- **Model 4354R disk enclosure**—Ultra3 SCSI with 14 drive bays, dual-bus I/O module. [Table 7](#) shows the addresses for each device in a three-shelf, dual-bus configuration. A maximum of three Model 4354R disk enclosures can be used with each Model 2200 controller enclosure.

Note: Appendix A contains storage set profiles you can copy and use to create your own system profiles. It also contains an enclosure template you can use to help you keep track of the location of devices and storage sets in your shelves.

Table 4: PTL addressing, single-bus configuration, six Model 4310R enclosures

Model 4310R Disk Enclosure Shelf 6 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200
Model 4310R Disk Enclosure Shelf 5 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200

Model 4310R Disk Enclosure Shelf 4 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200
Model 4310R Disk Enclosure Shelf 1 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk41000	Disk41100	Disk41200
Model 4310R Disk Enclosure Shelf 2 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk11000	Disk11100	Disk11200
Model 4310R Disk Enclosure Shelf 3 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk31000	Disk31100	Disk31200

Table 5: PTL addressing, dual-bus configuration, three Model 4350R enclosures

Model 4350R Disk Enclosure Shelf 1 (Single-bus)										
Bay	SCSI Bus A					SCSI Bus B				
	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400
Model 4350R Disk Enclosure Shelf 2 (Single-bus)										
Bay	SCSI Bus A					SCSI Bus B				
	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400
Model 4350R Disk Enclosure Shelf 3 (Single-bus)										
Bay	SCSI Bus A					SCSI Bus B				
	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400

Table 6: PTL addressing, single-bus configuration, six Model 4314R enclosures

Model 4314R Disk Enclosure Shelf 6 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk60900	Disk61000	Disk61100	Disk61200	Disk61300	Disk61400	Disk61500

Model 4314R Disk Enclosure Shelf 5 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk50900	Disk51000	Disk51100	Disk51200	Disk51300	Disk51400	Disk51500

Model 4314R Disk Enclosure Shelf 4 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk51100	Disk51200	Disk51300	Disk51400	Disk51500

Model 4314R Disk Enclosure Shelf 1 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500

Model 4314R Disk Enclosure Shelf 2 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500

Model 4314R Disk Enclosure Shelf 3 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

Table 7: PTL addressing, dual-bus configuration, three Model 4354A enclosures.

Model 4354R Disk Enclosure Shelf 1 (Dual-bus)														
Bay	SCSI Bus A							SCSI Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800
Model 4354R Disk Enclosure Shelf 2 (Dual-bus)														
Bay	SCSI Bus A							SCSI Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800
Model 4354R Disk Enclosure Shelf 3 (Dual-bus)														
Bay	SCSI Bus A							SCSI Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

Choosing a Container Type

Different applications may have different storage requirements. You probably want to configure more than one kind of container within your subsystem.

In choosing a container, you choose between independent disks (JBODs) or one of several storageset types, as shown in [Figure 23](#). The independent disks and the selected storageset may also be partitioned.

The storagesets implement RAID (Redundant Array of Independent Disks) technology. Consequently, they all share one important feature: each storageset, whether it contains two disk drives or ten, looks like one large, virtual disk drive to the host.

Containers




Figure 23: Storage container types

Table 8 compares the different kinds of containers to help you determine which ones satisfy your requirements.

Table 8: Comparison of container types

Container Name	Relative Availability	Request Rate (Read/Write) I/O Per Second	Transfer Rate (Read/Write) MB Per Second	Applications
Independent disk drives (JBOD)	Equal to number of JBOD disk drives	Comparable to single disk drive	Comparable to single disk drive	—
Stripeset (RAID 0)	Proportionate to number of disk drives; worse than single disk drive	Excellent if used with large chunk size	Excellent if used with small chunk size	High performance for non-critical data
Mirrorset (RAID1)	Excellent	Good/Fair	Good/Fair	System drives; critical files
RAIDset (RAID 3/5)	Excellent	Excellent/good	Read: excellent (if used with small chunk sizes) Write: good (if used with small chunk sizes)	High request rates, read-intensive, data lookup
Striped Mirrorset (RAID 0+1)	Excellent	Excellent if used with large chunk size	Excellent if used with small chunk size	Any critical response-time application

For a comprehensive discussion of RAID, refer to *The RAIDBOOK—A Source Book for Disk Array Technology*.

Creating a Storageset Profile

Creating a profile for your storagesets, partitions, and devices can simplify the configuration process. Filling out a storageset profile helps you choose the storagesets that best suit your needs and to make informed decisions about the switches you can enable for each storageset or storage device that you configure in your subsystem.

For an example of a storageset profile, see [Table 9](#). This table contains blank profiles that you can copy and use to record the details for your storagesets. Use the information in this chapter to help you make decisions when creating storageset profiles.

Table 9: Example of storageset profile

Type of Storageset:								
<input type="checkbox"/> Mirrorset	<input checked="" type="checkbox"/> RAIDset	<input type="checkbox"/> Stripeset	<input type="checkbox"/> Striped	<input type="checkbox"/> JBOD				
Mirrorset								
Storageset Name R1. Disk Drives D10300, D20300, D10400, D20400								
Unit Number D101								
Partitions:								
Unit #	Unit #	Unit #	Unit #	Unit #	Unit #	Unit #	Unit #	Unit #
%	%	%	%	%	%	%	%	%
RAIDset Switches:								
Reconstruction Policy			Reduced Membership			Replacement Policy		
<input checked="" type="checkbox"/> Normal (default)			<input checked="" type="checkbox"/> No (default)			<input checked="" type="checkbox"/> Best performance (default)		
<input type="checkbox"/> Fast			<input type="checkbox"/> Yes, missing:			<input type="checkbox"/> Best fit		
						<input type="checkbox"/> None		
Mirrorset Switches:								
Replacement Policy			Copy Policy			Read Source		
<input type="checkbox"/> Best performance (default)			<input type="checkbox"/> Normal (default)			<input type="checkbox"/> Least busy (default)		
<input type="checkbox"/> Best fit			<input type="checkbox"/> Fast			<input type="checkbox"/> Round robin		
<input type="checkbox"/> None						<input type="checkbox"/> Disk drive:		

Initialize Switches:

Chunk size	Save Configuration	Metadata
<input checked="" type="checkbox"/> Automatic (default)	<input type="checkbox"/> No (default)	<input checked="" type="checkbox"/> Destroy (default)
<input type="checkbox"/> 64 blocks	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> Retain
<input type="checkbox"/> 128 blocks		
<input type="checkbox"/> 256 blocks		

Unit Switches:

Caching	Access by following hosts enabled
Read caching <input checked="" type="checkbox"/>	<u>ALL</u>
Read-ahead caching	<u> </u>
Write-back caching <input checked="" type="checkbox"/>	<u> </u>
Write-through caching	<u> </u>

Planning Considerations for Storageset

This section contains the guidelines for choosing the storageset type needed for your subsystem:

- "Stripeset Planning Considerations", page 80
- "Mirrorset Planning Considerations", page 82
- "RAIDset Planning Considerations", page 84
- "Striped Mirrorset Planning Considerations", page 86
- "Storageset Expansion Considerations", page 88
- "Partition Planning Considerations", page 88

Stripeset Planning Considerations

Stripesets (RAID 0) enhance I/O performance by spreading the data across multiple disk drives. Each I/O request is broken into small segments called “chunks.” These chunks are then simultaneously “striped” across the disk drives in the storageset, thereby enabling several disk drives to participate in one I/O request.

For example, in a three-member stripeset that contains disk drives Disk 10000, Disk 20000, and Disk 10100, the first chunk of an I/O request is written to Disk 10000, the second to Disk 20000, the third to Disk 10100, the fourth to Disk 10000, until all of the data has been written to the drives (Figure 24).




Figure 24: Three-member RAID 0 stripeset (example 1)

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rate. You can set the chunk size or use the default setting (see "Chunk Size", page 93, for information about setting the chunk size). [Figure 25](#) shows another example of a three-member RAID 0 stripeset.

A major benefit of striping is that it balances the I/O load across all of the disk drives in the storageset. This can increase the subsystem performance by eliminating the hot spots (high localities of reference) that occur when frequently accessed data becomes concentrated on a single disk drive.




Figure 25: Three-member RAID 0 stripeset (example 2)

Keep the following points in mind as you plan your stripesets:

- Reporting methods and size limitations prevent certain operating systems from working with large stripesets.
- A storageset should only contain disk drives of the same capacity. The controller limits the effective capacity of each member to the capacity of the smallest member in the storageset (base member size) when the storageset is initialized. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you waste 5 GB of capacity on each 9 GB member.

If you need high performance and high availability, consider using a RAIDset, striped-mirrorset, or a host-based shadow of a stripeset.

- Striping does not protect against data loss. In fact, because the failure of one member is equivalent to the failure of the entire stripeset, the likelihood of losing data is higher for a stripeset than for a single disk drive.

For example, if the mean time between failures (MTBF) for a single disk is 1 hour, then the MTBF for a stripeset that comprises N such disks is $1/N$ hours. As another example, if the MTBF of a single disk is 150,000 hours (about 17 years), a stripeset comprising four of these disks would only have an MTBF of slightly more than 4 years.


For this reason, you should avoid using a stripeset to store critical data. Stripesets are more suitable for storing data that can be reproduced easily or whose loss does not prevent the system from supporting its critical mission.

- Evenly distribute the members across the device ports to balance the load and provide multiple paths.
- Stripesets may contain between two and 24 members.
- If you plan to use mirror members to replace failing drives, then create the original stripeset as a stripeset of 1-member mirrorsets.
- Stripesets are well-suited for the following applications:
 - Storing program image libraries or run-time libraries for rapid loading.
 - Storing large tables or other structures of read-only data for rapid application access.
 - Collecting data from external sources at very high data transfer rates.
- Stripesets are not well-suited for the following applications:
 - A storage solution for data that cannot be easily reproduced or for data that must be available for system operation.
 - Applications that make requests for small amounts of sequentially located data.
 - Applications that make synchronous random requests for small amounts of data.


Spread the member drives as evenly as possible across the six I/O device ports.

Mirrorset Planning Considerations

Mirrorsets (RAID 1) use redundancy to ensure availability, as illustrated in [Figure 26](#). For each primary disk drive, there is at least one mirror disk drive. Thus, if a primary disk drive fails, its mirror drive immediately provides an exact copy of the data. [Figure 27](#) shows a second example of a mirrorset.



CXO7288A

Figure 26: Mirrorsets maintain two copies of the same data**Figure 27: Mirrorset example 2**

Keep these points in mind when planning mirrorsets:

- Data availability with a mirrorset is excellent but comes with a higher cost—you need twice as many disk drives to satisfy a given capacity requirement. If availability is your top priority, consider using dual-redundant controllers and redundant power supplies.

- You can configure up to a maximum of 20 RAID 3/5 mirrorsets per controller or pair of dual-redundant controllers. Each mirrorset may contain up to 6 members. Refer to "[Configuration Rules for the Controller](#)", page 65, for detailed information on maximum numbers. 30 RAID 3/5 and RAID 1 mirrorsets are permitted, however, there is limit of no more than 20 RAID 3/5 mirrorsets in such a configuration.
- Both write-back cache modules must be the same size.
- A mirrorset should only contain disk drives of the same capacity.
- Spread mirrorset members across different device ports (drive bays).
- Mirrorsets are well-suited for the following:
 - Any data for which reliability requirements are extremely high
 - Data to which high-performance access is required
 - Applications for which cost is a secondary issue
- Mirrorsets are not well-suited for the following applications:
 - Write-intensive applications (a performance hit of 10 percent occurs)
 - Applications for which cost is a primary issue

RAIDset Planning Considerations

RAIDsets (RAID 3/5) are enhanced stripesets—they use striping to increase I/O performance and distributed-parity data to ensure data availability. [Figure 28](#) shows an example of a RAIDset that uses five members.




Figure 28: Five-member RAIDset using parity

RAIDsets are similar to stripesets in that the I/O requests are broken into smaller “chunks” and striped across the disk drives. RAIDsets also create chunks of parity data and stripe them across all the members of the RAIDset. Parity data is derived mathematically from the I/O data and enables the controller to reconstruct the I/O data if a single disk drive fails. Thus, it becomes possible to lose a disk drive without losing access to the data it contained. Data could be lost if a second disk drive fails before the controller replaces the first failed disk drive and reconstructs the data.

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rates. You can set the chunk size or use the default setting. See "["Chunk Size"](#)", page 93, for information about setting the chunk size.

Keep these points in mind when planning RAIDsets:

- Reporting methods and size limitations prevent certain operating systems from working with large RAIDsets.
- Both cache modules must be the same size.

- A RAIDset must include at least 3 disk drives, but no more than 14.
- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you waste 5 GB of capacity on each 9 GB member.
- RAIDsets are particularly well-suited for the following:
 - Small to medium I/O requests
 - Applications requiring high availability
 - High read request rates
 - Inquiry-type transaction processing
- RAIDsets are not particularly well-suited for the following:
 - Write-intensive applications
 - Database applications in which fields are continually updated
 - Transaction processing

Striped Mirrorset Planning Considerations

Striped mirrorsets (RAID 0+1) are a configuration of stripesets whose members are also mirrorsets ([Figure 29](#)). Consequently, this kind of storageset combines the performance of striping with the reliability of mirroring. The result is a storageset with very high I/O performance and high data availability. [Figure 30](#) shows a second example of a striped mirrorset using six members.




Figure 29: Striped mirrorset (example 1)

The failure of a single disk drive has no effect on the ability of the storageset to deliver data to the host. Under normal circumstances, a single disk drive failure has very little effect on performance. Because striped mirrorsets do not require any more disk drives than mirrorsets, this storageset is an excellent choice for data that warrants mirroring.




Figure 30: Striped mirrorset (example 2)

Plan the mirrorset members, and plan the stripeset that contains them. Review the recommendations in "[Planning Considerations for Stageset](#)", page 80, and "[Mirrorset Planning Considerations](#)", page 82.

Stageset Expansion Considerations

Stageset Expansion allows for the joining of two of the same kind of storage containers by concatenating RAIDsets, stripesets, or individual disks, thereby forming a larger virtual disk which is presented as a single unit. The *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide* describes the CLI command, ADD CONCATSETS, which is used to perform concatenation.



Caution: Use the ADD CONCATSETS command only with host operating systems that support dynamic volume expansion. Use of this command could result in inaccessible data, if the operating system cannot handle one of its disks increasing in size.

Partition Planning Considerations

Use partitions to divide a container (stageset or individual disk drive) into smaller pieces, each of which can be presented to the host as its own storage unit. [Figure 31](#) shows the conceptual effects of partitioning a single-disk container.




Figure 31: One example of a partitioned single-disk unit

You can create up to eight partitions per stageset (disk drive, RAIDset, mirrorset, stripeset, or striped mirrorset). Each partition has its own unit number so that the host can send I/O requests to the partition just as it would to any

unpartitioned storage set or device. Partitions are separately addressable storage units; therefore, you can partition a single storage set to service more than one user group or application.

Defining a Partition

Partitions are expressed as a percentage of the storage set or single disk unit that contains them:

- Mirrorsets and single disk units—the controller allocates the largest whole number of blocks that are equal to or less than the percentage you specify.
- RAIDsets and stripesets—the controller allocates the largest whole number of stripes that are less than or equal to the percentage you specify.
 - Stripesets—the stripe size = chunk size × number of members.
 - RAIDsets—the stripe size = chunk size × (number of members minus 1)

An unpartitioned storage unit has more capacity than a partition that uses the whole unit because each partition requires a small amount of disk space for metadata.

Guidelines for Partitioning Storage Sets and Disk Drives

Keep these points in mind when planning partitions for storage sets and disks:

- Each storage set or disk drive may have up to eight partitions.
- In transparent failover mode, all partitions of a particular container must be on the same host port. Partitions cannot be split across host ports.
- In multiple-bus failover mode, all the partitions of a particular container must be on the same controller. Partitions cannot be split across controllers.
- Partitions cannot be combined into storage sets. For example, you cannot divide a disk drive into three partitions, then combine those partitions into a RAIDset.
- Just as with storage sets, you do not have to assign unit numbers to partitions until you are ready to use them.
- The Clone utility cannot be used with partitioned mirrorsets or partitioned stripesets. (See "[Creating Clones for Backup](#)", page 179, for details about cloning.)

Changing Characteristics Through Switches

CLI command switches allow you another level of command options. There are three types of switches that modify the storageset and unit characteristics:

- Storageset switches
- Initialization switches
- Unit switches

The following sections describe how to enable/modify switches. They also contain a description of the major CLI command switches.

Enabling Switches

If you use SWCC to configure the device or storageset, you can set switches from SWCC during the configuration process, and SWCC automatically applies them to the storageset or device. See the SWCC online help for information about using SWCC.

The configuration procedure found in [Chapter 5](#) of this guide indicates when and how to enable each switch if you use CLI commands to configure the storageset or device manually. The *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide* contains the details of the CLI commands and their switches.

Changing Switches

You can change the RAIDset, mirrorset, device, and unit switches at any time. You cannot change the initialize switches without destroying data on the storageset or device. These switches are integral to the formatting and can only be changed by re-initializing the storageset.



Caution: Initializing a storageset is similar to formatting a disk drive; all data is destroyed during this procedure.

Specifying Stageset and Partition Switches

The characteristics of a particular stageset can be set by specifying switches when the stageset is added to the controllers' configuration. Once a stageset has been added, the switches can be changed by using a SET command. Switches can be set for partitions and the following types of stagesets:

- RAIDset
- Mirrorset

Stripesets have no specific switches associated with their ADD and SET commands.

RAIDset Switches

Use the following types of switches to control how a RAIDset ensures data availability:

- Replacement policy
- Reconstruction policy
- Remove/replace policy

For details on the use of these switches, refer to SET RAIDSET and SET *RAIDset-name* commands in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide*.

Mirrorset Switches

Use the following switches to control how a mirrorset behaves to ensure data availability:

- Replacement policy
- Copy speed
- Read source
- Membership

For details on the use of these switches, refer to ADD MIRRORSET and SET *mirrorset-name* commands in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide*.

Partition Switches

The following switches are available when creating a partition:

- Size
- Geometry

For details on the use of these switches, refer to CREATE_PARTITION command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide*.

Specifying Initialization Switches

Initialization switches set characteristics for established storagesets before they are made into units. The following kinds of switches effect the format of a disk drive or storageset:

- Chunk Size (for stripesets and RAIDsets only)
- Save Configuration
- Destroy/NodeDestroy
- Geometry

Each of these switches is described in the following sections.

Note: After initializing the storageset or disk drive, you cannot change these switches without reinitializing the storageset or disk drive.

Chunk Size

With ACS software, a parameter for chunk size (chunksize=default or n) on some storagesets can be set. However, unit performance may be negatively impacted if a non-default value is selected as the chunksize.

If a non-default chunk size has been calculated, verify that the chunk size value is divisible by 8, with no remainder. If the value is not aligned with this rule, adjust the chunk size value upward until it divisible by 8, with no remainder.

Specify the chunk size of the data to be stored to control the stripesize used in RAIDsets and stripesets:

- CHUNKSIZE=DEFAULT lets the controller set the chunk size based on the number of disk drives (d) in a stripeset or RAIDset. If number of drives is less or equal to 9, then chunk size = 256. If the number of drives is greater than 9, then chunk size = 128.
- CHUNKSIZE=*n* lets you specify a chunk size in blocks. The relationship between chunk size and request size determines whether striping increases the request rate or the data-transfer rate.

Increasing the Request Rate

A large chunk size (relative to the average request size) increases the request rate by enabling multiple disk drives to respond to multiple requests. If one disk drive contains all of the data for one request, then the other disk drives in the storage set are available to handle other requests. Thus, separate I/O requests can be handled in parallel, which increases the request rate. This concept is shown in [Figure 32](#).




Figure 32: Large chunk size increases request rate

Large chunk sizes also tend to increase the performance of random reads and writes. HP recommends that you use a chunk size of 10 to 20 times the average request size, rounded to the closest prime number.

To calculate the chunk size that should be used for your subsystem, you must first analyze the types of requests that are being made to the subsystem:

- Many parallel I/Os that use a small area of disk should use a chunk size of 10 times the average transfer request rate.
- Random I/Os that are scattered over all the areas of the disks should use a chunk size of 20 times the average transfer request rate.
- If you do not know, use a chunk size of 15 times the average transfer request rate.

- If you have mostly sequential reads or writes (like those needed to work with large graphic files), make the chunk size for RAID 0 and RAID 0+1 a small number (for example: 67 sectors). For RAID 5, make the chunk size a relatively large number (for example: 253 sectors).

Table 10 shows a few examples of chunk size selection.

Table 10: Example chunk sizes

Transfer Size (KB)	Small Area of I/O Transfers	Unknown	Random Areas of I/O Transfers
2	41	59	79
4	79	113	163
8	157	239	317

Increasing Sequential Data Transfer Performance

RAID 0 and RAID 0+1 sets intended for high data transfer rates should use a relatively low chunk size (for example: 67 sectors). RAID 5 sets intended for high data rate performance should use a relatively large number (for example: 253 sectors).

Save Configuration

The SAVE_CONFIGURATION switch is for a *single-controller configuration only*. This switch reserves an area on each of the disks for the container being initialized. The controller can write subsystem configuration data on this area. If the controller is replaced, the new controller can read the subsystem configuration from the reserved areas of disks.

If you specify SAVE_CONFIGURATION for a multi-device storage set, such as a stripeset, the complete subsystem configuration is periodically written on each disk in the storage set.

The SHOW DEVICES FULL command shows which disks are used to backup configuration information.

Note: DO NOT use SAVE_CONFIGURATION in dual redundant controller installations. It is not supported and may result in unexpected controller behavior.

Note: HP recommends that you DO NOT use SAVE_CONFIGURATION on every unit and device on the controller.

Destroy/Nodestroy

Specify whether to destroy or retain your data and metadata when a disk is initialized after it has been used in a mirrorset or as a single-disk unit.

Note: The DESTROY and NODESTROY switches are only valid for mirrorsets and striped mirrorsets.

- DESTROY (default) overwrites your data and forced-error metadata when a disk drive is initialized.
- NODESTROY preserves your data and forced-error metadata when a disk drive is initialized. Use NODESTROY to create a single-disk unit from any disk drive that has been used as a member of a mirrorset. See the REDUCED command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide* for information on removing disk drives from a mirrorset.

NODESTROY is ignored for members of a RAIDset.

Geometry

The geometry parameters of a storageset can be specified. The geometry switches are:

- CAPACITY—The number of logical blocks. The range is from 1 to the maximum container size.
- CYLINDERS—The number of cylinders used. The range is from 1 to 16777215.
- HEADS—The number of disk heads used. The range is from 1 to 255.
- SECTORS_PER_TRACK—The number of sectors per track used. The range is from 1 to 255.

Specifying Unit Switches

Several switches control the characteristics of units. The unit switches are described under the `SET unit-number` command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Command Line Interface Reference Guide*.

One unit switch, `ENABLE/DISABLE_ACCESS_PATH`, determines which host connections can access the unit, and it is described in the larger topic of matching units to specific hosts. This topic is covered in "[Determining Connections](#)", page 41.

Creating Storage Maps

Configuring a subsystem is easier if you know how the storage sets, partitions, and JBODs correspond to the disk drives in your subsystem. You can more easily see this relationship by creating a hardcopy representation, also known as a storage map.

To make a storage map, fill out the templates provided in [Appendix A](#) as you add storage sets, partitions, and JBOD disks to the configuration and assign them unit numbers. Label each disk drive in the map with the higher levels it is associated with, up to the unit level.

Using LOCATE Command to Find Devices

If you want to complete a storage map at a later time but do not remember where the disk drives and partitions are located, use the CLI command LOCATE. The LOCATE command flashes the (fault) LED on the drives associated with the specific storage set or unit. To turn off the flashing LEDs, enter the CLI command LOCATE CANCEL.

To locate all the disk drives that make up unit D104:

1. Enter the following command:

```
LOCATE D104
```

The LEDs on the disk drives that make up unit D104 flashes.

2. Note the position of all the drives contained within D104.
3. Enter the following command to turn off the flashing LEDs:

```
LOCATE CANCEL
```

To locate all the drives that make up RAIDset R1:

1. Enter the following command:

```
LOCATE R1
```

2. Note the position of all the drives contained within R1.
3. Enter the following command to turn off the flashing LEDs:

```
LOCATE CANCEL
```

Example Storage Map—Model 4310R Disk Enclosure

[Table 11](#) shows an example of four Model 4310R disk enclosures (single-bus I/O).

- Unit D100 is a 4-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, Disk30000, and Disk40000.
- Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:
 - M1 is a 2-member mirrorset consisting of Disk10100 and Disk20100.
 - M2 is a 2-member mirrorset consisting of Disk30100 and Disk40100.
- Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10200 and Disk20200.
- Unit D103 is a 2-member mirrorset named M4. M4 consists of Disk30200 and Disk40200.
- Unit D104 is 3-member stripeset named S2. S2 consists of Disk10300, Disk20300, and Disk30300.
- Unit D105 is a single (JBOD) disk named Disk40300.
- Unit D106 is a 3-member RAID 3/5 storage set named R2. R2 consists of Disk10400, Disk20400, and Disk30400.
- Unit D107 is a single (JBOD) disk named Disk40400.
- Unit D108 is a 4-member stripeset named S3. S3 consists of Disk10500, Disk20500, Disk30500, and Disk40500.
- Unit D1 is a 2-member striped mirrorset named S4. S4 consists of M4 and M5:
 - M5 is a 2-member mirrorset consisting of Disk10800 and Disk20800.
 - M6 is a 2-member mirrorset consisting of Disk30800 and Disk40800.
- Unit D2 is a 4-member RAID 3/5 storage set named R3. R3 consists of Disk11000, Disk21000, Disk31000, and Disk41000.
- Unit D3 is a 4-member stripeset named S5. S5 consists of Disk11100, Disk21100, Disk31100, and Disk41100.
- Unit D4 is a 2-member mirrorset named M7. M7 consists of Disk11200 and Disk21200.
- Disk31200 and Disk41200 are spareset members.

Table 11: Model 4310R disk enclosure, example of storage map

Model 4310R Disk Enclosure Shelf 4 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1 Disk40000	D101 S1 M2 Disk40100	D103 M4 Disk40200	D105 Disk40300	D107 Disk40400	D108 S3 Disk40500	D1 S4 M6 Disk40800	D2 R3 Disk41000	D3 S5 Disk41100	spare Disk41200
Model 4310R Disk Enclosure Shelf 1 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1 Disk10000	D101 S1 M1 Disk10100	D102 M3 Disk10200	D104 S2 Disk10300	D106 R2 Disk10400	D108 S3 Disk10500	D1 S4 M5 Disk10800	D2 R3 Disk11000	D3 S5 Disk11100	D4 M7 Disk11200
Model 4310R Disk Enclosure Shelf 2 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1 Disk20000	D101 S1 M1 Disk20100	D102 M3 Disk20200	D104 S2 Disk20300	D106 R2 Disk20400	D108 S3 Disk20500	D1 S4 M5 Disk20800	D2 R3 Disk21000	D3 S5 Disk21100	D4 M7 Disk21200

Model 4310R Disk Enclosure Shelf 3 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1	D101 S1 M2	D103 M4	D104 S2	D106 R2	D108 S3	D1 S4 M6	D2 R3	D3 S5	spare
Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk31000	Disk31100	Disk31200	

3

Preparing the Host System

This chapter describes how to prepare your IBM AIX host computer to accommodate the HSG80 controller storage subsystem.

The following information is included in this chapter:

- ["Installing RAID Array Storage System"](#), page 104
- ["Making a Physical Connection"](#), page 108
- ["HP StorageWorks Solution Software Installation"](#), page 110

Installing RAID Array Storage System



WARNING: A shock hazard exists at the backplane when the controller enclosure bays or cache module bays are empty.

Be sure the enclosures are empty, then mount the enclosures into the rack. DO NOT use the disk enclosure handles to lift the enclosure. The handles cannot support the weight of the enclosure. Only use these handles to position the enclosure in the mounting brackets.

Use two people to lift, align, and install any enclosure into a rack. Failure to use two people might cause personal injury and/or equipment damage.



Caution: Controller and disk enclosures have no power switches. Make sure the controller enclosures and disk enclosures are physically configured before turning the PDU on and connecting the power cords. Failure to do so can cause equipment damage.

1. Be sure the enclosures are empty before mounting them into the rack. If necessary, remove the following elements from the controller enclosure:
 - Environmental Monitoring Unit (EMU)
 - Power Supplies
 - External Cache Batteries (ECBs)
 - Fans
2. If necessary, remove the following elements from the disk enclosure:
 - Power Supply/Blower Assemblies
 - Disk Drives
 - Environmental Monitoring Unit (EMU)
 - I/O Modules
3. Refer to the *HP StorageWorks Model 2100 and 2200 Ultra SCSI Controller Enclosures User Guide* and the *HP StorageWorks Model 4300 Family Ultra3 LVD Disk Enclosures User Guide* for further information.
4. Install brackets onto the controller enclosure and disk enclosures. Using two people, mount the enclosures into the rack. Refer to the mounting kit documentation for further information.

5. Install the elements. Install the disk drives. Make sure you install blank panels in any unused bays.

Fibre Channel cabling information is shown to illustrate supported configurations. In a dual-bus disk enclosure configuration, disk enclosures 1, 2, and 3 are stacked below the controller enclosure—two SCSI Buses per enclosure (see [Figure 33](#)). In a single-bus disk enclosure configuration, disk enclosures 6, 5, and 4 are stacked above the controller enclosure and disk enclosures 1, 2, and 3 are stacked below the controller enclosure—one SCSI Bus per enclosure (see [Figure 34](#)).

6. Connect the six VHDCI UltraSCSI bus cables between the controller and disk enclosures, as shown in [Figure 33](#) for a dual-bus system and [Figure 34](#) for a single-bus system. Note that the supported cable lengths are 1, 2, 3, 5, and 10 meters.
7. Connect the AC power cords from the appropriate rack AC outlets to the controller and disk enclosures.





Figure 33: Dual-bus enterprise storage RAID array storage system



CXO7382A

- | | | | |
|---|------------------|---|---------------------|
| ① | SCSI Bus 1 Cable | ② | SCSI Bus 2 Cable |
| ③ | SCSI Bus 3 Cable | ④ | SCSI Bus 4 Cable |
| ⑤ | SCSI Bus 5 Cable | ⑥ | SCSI Bus 6 Cable |
| ⑦ | AC Power Inputs | ⑧ | Fibre Channel Ports |

Figure 34: Single-bus enterprise storage RAID array storage system

Making a Physical Connection

To attach a host computer to the storage subsystem, install one or more host bus adapters into the computer. A Fibre Channel (FC) cable goes from the host bus adapter to an FC switch or hub.

Preparing to Install Host Bus Adapter

Before installing the host bus adapter, perform the following steps:

1. Perform a complete backup of the entire system.
2. Shut down the computer system or perform a hot addition of the adapter based upon directions for that server.

Installing Host Bus Adapter

To make a physical connection, first install a host bus adapter.



Caution: Protect the host bus adapter board from electrostatic discharge by wearing an ESD wrist strap. DO NOT remove the board from the antistatic cover until you are ready to install it.

You need the following items to begin:

- Host bus adapter board and the installation instructions included with the adapter
- The HP StorageWorks software installation kit
- The computer hardware manual
- Appropriate tools to service your computer

The host bus adapter board plugs into a standard PCI slot in the host computer. Refer to the system manual for instructions on installing PCI devices.

If installing the RAID subsystem on an High Availability Cluster Multi-Processor (HACMP) install a Fibre Quick Fibre Channel Host Bus Adapter (HBA) into each IBM host server in the cluster.

Note: Take note of the worldwide name (WWN) of each adapter. Do not power on anything yet. For the FC switches to autoconfigure, power on equipment in a certain sequence. Also, the controllers in the subsystem are not yet configured for compatibility with IBM AIX.

The Host Bus Adapter (HBA) installation is accomplished by first installing the HBA driver. The HBA driver is installed as part of the Solution Software installation.

Note: If you are installing the RAID subsystem on an HACMP cluster, install the HBA and its driver into each IBM host server in the cluster. For further information on HACMP, see [Appendix C](#).

HP StorageWorks Solution Software Installation

There are two types of HP StorageWorks solution kit installation procedures. You can perform a new installation or an upgrade installation.

- Perform a *new* installation if your system does not have a Cambex driver or any HP StorageWorks software installed.
- Perform an *upgrade* installation if you currently have a Cambex driver and/or HP StorageWorks software installed.

New Installation

This type of installation assumes that you have no pre-existing solution software for HSG80 installed. To install the solution software, perform the following steps:

1. Create a directory on the target IBM server where the installation kit will reside by typing the command:
`mkdir /AIX_Install`
2. Obtain the IBM AIX Solution Kit for HSG80 from the website
<http://h18004.www1.hp.com/products/storageworks/ma8kema12k/kits.html>.
3. Click on the link for the IBM AIX Solution Kit.
4. Click the **Download** button and save the `tar` file to the install directory you created on the target system.
5. Change the directory to the AIX install path by typing the command:
`cd AIX_Install`

6. Extract the kit from the `tar` file by typing the command:
`tar xvf <AIX tar file>`

7. Execute the installation script by issuing the command:
`./INSTALL`

The following is the output of the installation script. Notes and explanations have been added.

```
# ./INSTALL
```

First, the Cambex adapter driver, V1.5.26.0, is installed. The name of the driver fileset is PC1000.driver.obj

```
Solution Software Installation for AIX V88  
-----  
Installing Cambex adapter driver for AIX 4.3.3/5.1/5.2...OK  
Issuing bosboot command...OK
```

Next, entries to ODM database are added so that the OS supports both HSG60 and HSG80 controller-based subsystems. Since this is a new installation, the installation script looks for devices by running the cfgmgr command. The fileset names for the ODM entries are hsg60.rte and hsg80.rte respectively

```
Device Support Filesets  
-----  
Installing HSG60 ODM entries...OK  
Installing HSG80 ODM entries...OK  
Issuing bosboot Command... OK  
Searching for HSG80 Devices...OK
```

The HSG80 devices that are found by the script are listed below.

Note: This output assumes that there are already units defined on the HSG80 controllers for this server.

```
HSG80 Devices  
-----  
hdisk3 Available 20-60-00-0,0 DEC HSG80 Command Console LUN  
hdisk4 Available 20-60-00-0,1 DEC HSG80 RAID Array  
hdisk5 Available 20-60-00-0,2 DEC HSG80 RAID Array  
hdisk6 Available 20-60-00-0,3 DEC HSG80 RAID Array  
hdisk7 Available 20-60-00-1,0 DEC HSG80 Command Console LUN  
hdisk8 Available 20-60-00-1,1 DEC HSG80 RAID Array  
hdisk9 Available 20-60-00-1,2 DEC HSG80 RAID Array  
hdisk10 Available 20-60-00-1,3 DEC HSG80 RAID Array  
  
---Press ENTER to Continue---
```

You are now presented with the choice to install the SWCC Agent package. The SWCC Agent is a monitoring utility that provides email and error log notification if there is an error on a monitored RAID array. The SWCC Agent also allows the SWCC Client to monitor and configure HSG80 RAID Arrays. The SWCC Agent does not have to be installed but it is recommended. The Agent fileset is called swcc.rte.

SWCC Agent Package

```
Would you like to install/upgrade the SWCC Agent Now (Y/n)? Y  
Installing Agent Package...OK
```

The next part of the installation script allows you to configure Client systems that connect to the SWCC Agent, RAID subsystems that are monitored, and email addresses that are notified in the event of a subsystem error. The configuration can be done here or at a later date using the /usr/stgwks2/config.sh script or SMIT (**Applications->Compaq SWCC Agent Configuration**).

```
Would you like to configure the agent now (y/N)? y  
Copyright 2002 Hewlett-Packard Corporation.  
Copyright (c) Compaq Computer Corporation, 2002  
All Rights Reserved. Unpublished rights reserved under the  
copyright laws of the United States.
```

The software contained on this media is proprietary to and embodies the confidential technology of Compaq Computer Corporation. Possession, use, duplication or dissemination of the software and media is authorized only pursuant to a valid written license from Compaq Computer Corporation.

RESTRICTED RIGHTS LEGEND Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in Subparagraph (c)(1)(ii) of DFARS 252.227-7013, or in FAR 52.227-19, as applicable.

```
Enterprise Storage Management (steam)  
Agent Configuration Script  
You must be root to use this script!  
Current user is root.  
Do you wish to continue? [Y,n]:
```

At this point you are prompted to add a Client system. The Client system must be an Intel 32-bit based machine running either Windows NT 4.0, Windows 2000, or Windows Server 2003 and it must use a fixed IP address that is resolvable through DNS or NIS. Here is the guide to the meaning of the communication and monitoring options.

The Access Privilege options are:

- 0—The Client can only view the main storage window; it cannot open a system window for any of the RAID arrays on the host that you are currently on
- 1—The Client can view the main storage window and any of the system windows. However, this Client cannot do any configuration.
- 2—This Client can view any storage window and configure any RAID array on the current host system

The Notification Options are:

- 0—There is no asynchronous error notification to the Client.
- 1—The Client machine receives notification via TCP/IP only. Only the SWCC Client receives error messages.
- 2—The Client machine is notified by SMNP only. This means that errors sent out by the Agent can only be received by an SMNP Client. The SWCC Client does not report any errors.
- 3—The Client machine is notified by TCP/IP and SMNP. Both the SWCC Client and an SMNP monitoring Client can receive errors messages from the Agent.

In the following example we are adding a Client called client1 with full monitoring privileges and error notification to client1 via TCP/IP and SNMP traps

Adding a Client system

Enter the host name of the Client system: client1

Enter the subsystem access privilege (which controls the level of access a client has for communicating with the storage subsystems).

The possible options are:

```
0 = Navigation Tree monitoring Only  
1 = Navigation Tree and Storage Window monitoring Only  
2 = Navigation Tree monitoring, Storage Window monitoring and Configuration
```

Enter Access Level (0, 1, 2): 2

The Agent server can notify a client when an error condition occurs. Notification schemes available are:

```
0 = No Error Notification  
1 = Notification via a TCP/IP Socket  
2 = Notification via the SNMP protocol  
3 = Notification via both TCP/IP and SNMP
```

Enter Error Notification Level (0, 1, 2, 3): 3

Adding client--

```
name: client1  
access level: 2  
error notification: 3
```

Is this information correct? [Y,n]

Would you like to add another client? [y,N]

Here you can add a storage subsystem to monitor. You can choose any device as an access device, however HP StorageWorks recommends that you use the Command Console LUN devices. It is also recommended that you accept the default monitoring level of 300 seconds. The monitoring interval is the amount of time between polls of the monitored subsystems.

```
Adding storage subsystems
```

```
Command Console Access Devices
```

```
-
```

name	status	location	description
hdisk4	Available	20-60-00-0,1	DEC HSG80 RAID Array
hdisk5	Available	20-60-00-0,2	DEC HSG80 RAID Array
hdisk6	Available	20-60-00-0,3	DEC HSG80 RAID Array
hdisk8	Available	20-60-00-1,1	DEC HSG80 RAID Array
hdisk9	Available	20-60-00-1,2	DEC HSG80 RAID Array
hdisk10	Available	20-60-00-1,3	DEC HSG80 RAID Array
hdisk3	Available	20-60-00-0,0	DEC HSG80 Command Console LUN
hdisk7	Available	20-60-00-1,0	DEC HSG80 Command Console LUN

```
Specify access device to query for information ("hdisk5"):  
hdisk3
```

```
The following subsystem has been found:
```

```
Controller type: HSG80
```

```
Firmware version: YC04
```

```
Serial numbers: ZG11800073, ZG10707148
```

```
EMU type : None installed
```

```
Would you like to add this subsystem to the 'storage.ini' file?
```

```
Enter a name for this subsystem: raid1
```

```
Enter a monitoring interval, in secs
```

```
(0 => no monitoring, default is 300):
```

```
Would you like to add another subsystem? [y,N]
```

Next, add users and/or email addresses to notify in the event of a subsystem error.

The different notification levels are:

1—Only notification for serious errors that need immediate attention is sent.

2—Notification for serious errors and warning messages is sent.

3—Notification for serious errors, warning messages, and informational messages is sent.

```
Adding a User for Mail Notification
```

```
-----
```

```
Enter the email address of the user: root
```

```
Enter the error notification level for this user.
```

```
The user will be notified of errors at this level and above.
```

```
The possible options are:
```

```
1 = Fatal Errors
```

```
2 = Warning and Fatal Errors
```

```
3 = Info, Warning and Fatal Errors
```

```
Enter Notification Level (1, 2, 3): 2
```

```
Adding Notification User-- name: root
```

```
Error Notification Level: Warning and Fatal Errors
```

```
Is this information correct? [Y,n]
```

```
Would you like to add another user? [y,N]
```

```
Authorized mail notification list
```

```
-----
```

```
Username Notification Level
```

```
-----
```

```
root Warning & Fatal Errors
```

```
-- Hit RETURN to continue
```

```
The STEAM Agent package has been successfully configured.
```

Here you are given the opportunity to start the SWCC Agent. If you do not want to start the SWCC Agent, then you can start it later using the /usr/stgwks2/config.sh script or SMIT (**Applications->Compaq SWCC Agent Configuration**).

Would you like to execute the Agent now? [Y,n]

The STEAM agent has been started. The configuration process is complete.

-- Hit RETURN to continue --

The final part of the installation is the installation of the HSG80 Snapshot fileset. This allows you to create a Snapshot on the HSG80 of a particular disk in the OS. You must have the “S” variant of the V8.8-1 firmware to use this feature. A description of how to create Snapshots is included later in this document. The name of this fileset is hsg_snapshot.rte.

HSG80 Snapshots

HSG80 Snapshots allow you to make snapshots of specific disks on a RA8000, MA8000, or EMA12000

Would you like to install the HSG80 snapshot fileset? (Y/n)
Installing HSG80 Snapshot fileset...OK

At this point the installation is complete. The last line states that you can uninstall the solution software by using the cpq_hsg_uninstall command or using the UNINSTALL command in the installation directory.

The solution software can be uninstalled anytime by issuing the cpq_hsg_uninstall command or running the UNINSTALL command in the installation directory

Installation Complete.

Upgrading Solution Software

To upgrade the Solution Software:

- Upgrade the Cambex adapter driver.
- Upgrade the HSG60 and/or HSG80 ODM entries.
- Upgrade the build of the SWCC Agent.

To upgrade the solution software:

1. Create a directory on the target IBM server where the installation kit will reside by typing the command:
`mkdir /AIX_Install`
2. Obtain the IBM AIX Solution Kit for HSG80 from the website
<http://h18004.www1.hp.com/products/storageworks/ma8kema12k/kits.html>.
3. Click on the link for the IBM AIX Solution Kit.
4. Click the **Download** button and save the `tar` file to the install directory you created on the target system.
5. Change the directory to the AIX install path by typing the command:

```
cd AIX_Install
```

6. Extract the kit from the `tar` file by typing the command:
`tar xvf <AIX tar file>`
7. Execute the installation script by issuing the command:
`./INSTALL`

The following is the output of the installation script. Notes and explanations have been added.

```
# ./INSTALL
```

First, the installation script checks for a previous version of the Cambex adapter driver. If it finds one, it asks if you want to upgrade the driver. There is a driver upgrade script that is included with this solution kit. If you do not want to upgrade the driver automatically, you need to:

- Unmount any filesystems using HSG80 or HSG60 devices.
- Export any volume groups using any HSG80 or HSG60 devices.
- Remove any child devices of the Cambex adapter devices.
- Remove any Cambex adapter devices.
- Remove the Cambex adapter driver (`installp -u PC1000.driver.obj`).
- Install the new driver from the AIX installation path
(`installp -acd/ AIX_Install/driver/PC1000.image all`).
- Search for devices (`cfgmgr`).
- Import any previously exported volume groups.
- Mount any previously mounted filesystems to upgrade the driver.

Removing the Solution Software

To uninstall the solution software:

1. Stop all processes accessing any HSG80 or HSG60 devices.
2. Remove any volume groups on any HSG80 or HSG60 devices.
3. Issue the `cpq_hsg_uninstall` command or run the UNINSTALL program from the AIX installation kit. The following is the output from the `cpq_hsg_uninstall` command or UNINSTALL program.

```
# cpq_hsg_uninstall
This script will uninstall the following filesets:
hsg_snapshot.rte: create disk snapshots
swcc.rte: SWCC agent
hsg80.rte: HSG80 device support
hsg60.rte: HSG60 devices support
PC1000.driver.obj (optional): Cambex adapter driver
Before you execute this uninstall, be sure that any volume
groups that are attached to any HSG80/HSG60 disks are removed!
Do you want to proceed with uninstalling these filesets (Y/n)?
Removing hsg_snapshot.rte...OK
Removing swcc.rte...OK
Removing HSG80 devices...OK
Removing hsg80.rte...OK
Removing HSG60 devices...OK
Removing hsg60.rte...OK
```

Note: If you are using any other storage, such as the Enterprise Virtual Array, you should not uninstall the Cambex adapter driver.

```
Do you want to uninstall the Cambex driver (Y/n)?
Removing Cambex adapters...OK
Removing Cambex adapter driver...OK
Uninstall Complete.
```

Solution Software Upgrade Procedures

Use the following procedures for upgrades to your Solution Software. It is considered best practice to follow this order of procedures:

1. Perform backups of data prior to upgrade.
2. Verify operating system versions, upgrade operating systems to supported versions and patch levels.
3. Quiesce all I/O and unmount all file systems before proceeding.
4. Upgrade switch firmware.
5. If installing an operating system that uses Secure Path, upgrade Secure Path to the latest version which will also upgrade Solution Software, then skip to step 7.
6. Upgrade Solution Software.
7. Upgrade ACS software.

Note: Solely for the purpose of performing upgrades to the ACS firmware, this Solution Software Kit supports previous ACS V8.6 and 8.7. It is recommended that you do not mix ACS versions in the same SAN.

Refer to the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Maintenance and Service Guide* and the *Solution Software Release Notes* for the latest information on upgrades.

4

Installing and Configuring HSG Agent

StorageWorks Command Console (SWCC) enables real-time configuration of the storage environment and permits you to monitor and configure the storage connected to the HSG80 controller.

The following information is included in this chapter:

- "Why Use StorageWorks Command Console (SWCC)?", page 124
- "Installation and Configuration Overview", page 126
- "About the Network Connection for the Agent", page 127
- "Before Installing the Agent", page 129
- "Installing the Agent", page 129
- "Configuration Options for the SWCC Agent", page 130

Refer to [Chapter 5](#) for a description of how to configure a subsystem that uses Fibre Channel fabric topology.

Why Use StorageWorks Command Console (SWCC)?

StorageWorks Command Console (SWCC) enables you to monitor and configure the storage connected to the HSG80 controller. SWCC consists of Client and Agent.

- The Client provides pager notification and lets you manage your virtual disks. The Client runs on Windows 2000 with Service Pack 4, Windows NT 4.0 with Service Pack 6A or above, and Windows Server 2003 (32-bit).
- The Agent obtains the status of the storage connected to the controller. It also passes the status of the devices connected to the controller to other computers and provides email notification and error logging.

To receive information about the devices connected to your HSG80 controller over a TCP/IP network, you must install the Agent on a computer that is connected to a controller.

The Agent can also be used as a standalone application without Client. In this mode, which is referred to as Agent only, Agent monitors the status of the subsystem and provides local and remote notification in the event of a failure. A subsystem includes the HSG80 controller and its devices. Remote and local notification can be made by email, System Error Log and/or SNMP messages to an SNMP monitoring program.

Table 12: SWCC features and components

Features	Agent Required?	Client Required?
Creation of RAID sets: Striped device group (RAID 0) Mirrored device group (RAID 1) Striped mirrored device group (RAID 0+1) Striped parity device group (RAID 3/5) Individual device (JBOD)	Yes	Yes
Monitor multiple subsystems at once	Yes	No
Event logging	Yes	No
Email notification	Yes	No
Pager notification	Yes	Yes

Note: For serial and SCSI connections, the Agent is not required for creating virtual disks.

Installation and Configuration Overview

[Table 13](#) provides an overview of the installation.

Table 13: Installation and configuration overview


Step	Procedure
1	Verify that your hardware has been set up correctly. See the previous chapters in this guide.
2	Verify that you have a network connection for the Client and Agent systems. See " About the Network Connection for the Agent ", page 127.
3	Verify that there is a LUN for communications. This can be either the CCL or a LUN that was created with the CLI. See " What is the Command Console LUN? ", page 39, in Chapter 1 .
4	Install the Agent (TCP/IP network connections) on a system connected to the HSG80 controller.
5	Add the name of the Client system to the Agent's list of Client system entries (TCP/IP network connections). This can be done during installation or when reconfiguring the Agent.
6	Install the Client software on Windows 2000 with Service Pack 4 or Windows NT 4.0 with Service Pack 6A, and Windows Server 2003 (32-bit). See Appendix B .
7	Add the name of the Agent system to the Navigation Tree of each Client system that is on the Agent's list of Client system entries (TCP/IP network connections). See Appendix B .
8	Set up pager notification (TCP/IP network connections). Refer to "Setting Up Pager Notification" in the <i>HP StorageWorks Command Console V2.5 User Guide</i> .

About the Network Connection for the Agent

The network connection, shown in [Figure 35](#), displays the subsystem connected to a hub or a switch. SWCC can consist of any number of Clients and Agents in a network. However, it is suggested that you install only one Agent on a computer.

By using a network connection, you can configure and monitor the subsystem from anywhere on the LAN. If you have a WAN or a connection to the Internet, monitor the subsystem with TCP/IP.

Note: SWCC does not support the dynamic host configuration protocol (DHCP) or the Windows Internet Name Service (WINS).



CXO7240A

- | | |
|---|---|
| ① Agent system (has the Agent software) | ⑤ Hub or switch |
| ② TCP/IP network | ⑥ HSG80 controller and its device subsystem |
| ③ Client system (has the Client software) | ⑦ Servers |
| ④ Fibre Channel cable | |

Figure 35: An example of a network connection

Before Installing the Agent

The Agent requires the minimum system requirements, as defined in the release notes for your operating system. The program is designed to operate with the Client V2.5 on Windows 2000, Windows NT, and Windows Server 2003.

Installing the Agent

The SWCC Agent is installed as part of the Solution Software Installation. For details about installing the SWCC Agent, see [Chapter 3](#).

Event Log File

The Agent logs any subsystem events in `/var/adm/ras/steamd.log`.

Configuration Options for the SWCC Agent

After the Agent has been installed, it is possible to configure the SWCC Agent to change some of its operating settings or to add a new Client, subsystem, or email address. The further configuration options can be done from one of two places:

- The `/usr/stgwks2/config.sh` script
- SMIT (under **Applications->Compaq SWCC Agent Configuration**)

Note: If any change is made the SWCC Agent configuration, the Agent has to be stopped and restarted for the change to take effect.

All these options should be available from both places, except for password changes, which have to be done from the `config.sh` script.

- Change Enterprise OID—Change the Enterprise Object ID that is used when the Agent sends out SNMP traps.
- Change Password—Change the password that is used for authentication in the Client. You can only change the password using the `/usr/stgwks2/config.sh` script.
- Start/Stop Agent—Start or stop the Agent process
- Toggle Agent Startup on Boot—Set up whether the Agent starts automatically when the system is booted. Automatic start is done through an entry placed in the `/etc/inittab` file. When Agent startup on boot is enabled, the setting in `/etc/inittab` file is set to “once”. When Agent startup on boot is disabled, it is set to “off”.
- Add Client—Add another Client PC. There are different levels of access privilege that can be set and different methods of notification that can be set for each Client. Here is a brief description of each:

The Access Privilege options are:

- 0—The Client can only view the main storage window; it cannot open a system window for any of the RAID arrays on the host that you are currently on
- 1—The Client can view the main storage window and any of the system windows. However, this Client cannot do any configuration.
- 2—This Client can view any storage window and configure any RAID array on the current host system

The Notification Options are:

- 0—There is no asynchronous error notification to the Client.
- 1—The Client machine receives notification via TCP/IP only. Only the SWCC Client receives error messages.
- 2—The Client machine is notified by SMNP only. This means that errors sent out by the Agent can only be received by an SMNP Client. The SWCC Client does not report any errors.
- 3—The Client machine is notified by TCP/IP and SMNP. Both the SWCC Client and an SMNP monitoring Client can receive errors messages from the Agent.
- Remove Client—Remove a previously defined Client system.
- Show/Edit Client Attributes—Change the access privileges or notification options for a Client that has already been defined.
- Add RAID Subsystem—Add another RAID subsystem to monitor. The two pieces of information you need for this are which device you want to use as a communication devices (preferably a Command Console Device) and the monitoring interval (the duration in between subsystem polls).
- Remove RAID Subsystem—Remove a previously defined subsystem.
- Show/Edit RAID Subsystem Attributes—Edit a previously defined RAID subsystem and change its access device or monitoring interval.
- Toggle Syslog Notification—Enable or disable event notification to the /var/adm/ras/steamd.log file.
- Add User To Mail List—Add an email address that the Agent notifies if there is a subsystem event. The options for email notification are:
 - 1—Only notification for serious errors that need immediate attention are sent.
 - 2—Notification for serious errors and warning messages are sent.
 - 3—Notification for serious errors, warning messages, and informational messages are sent.
- Remove User From Mail List—Remove a previously defined email address.
- Show/Edit Mail user Attributes—Change the notification attribute of a previously defined email address.
- Toggle Mail Notification—Enable or disable email notification if there is a subsystem event.

FC Configuration Procedures

5

This chapter describes procedures to configure a subsystem that uses Fibre Channel (FC) fabric topology. In fabric topology, the controller connects to its hosts through hubs.

The following information is included in this chapter:

- "Establishing a Local Connection", page 134
- "Setting Up a Single Controller", page 135
- "Setting Up a Controller Pair", page 145
- "Configuring Devices", page 152
- "Configuring Storage Containers", page 153
- "Assigning Unit Numbers and Unit Qualifiers", page 159
- "Configuration Options", page 160
- "Verifying Storage Configuration from Host", page 164

Use the command line interpreter (CLI) or StorageWorks Command Console (SWCC) to configure the subsystem. This chapter uses CLI to connect to the controller. To use SWCC for configuration, see the SWCC online help for assistance.


Note: These configuration procedures assume that controllers and cache modules are installed in a fully functional and populated enclosure and that the PCMCIA cards are installed.

To install a controller or cache module and the PCMCIA card, see the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Maintenance and Service Guide*.

Establishing a Local Connection

A local connection is required to configure the controller until a command console LUN (CCL) is established using the CLI. Communication with the controller can be through the CLI or SWCC.

The maintenance port, shown in [Figure 36](#), provides a way to connect a maintenance port. The maintenance port can be an EIA-423 compatible terminal or a computer running a terminal emulator program. The maintenance port accepts a standard RS-232 jack. The maintenance port cable shown in [Figure 36](#) has a 9-pin connector molded onto the end for a PC connection. If you need a terminal connection or a 25-pin connection, you can order optional cabling.



❶ Maintenance Port
Cable ❷ Maintenance Port
Port

Figure 36: Maintenance port connection



Caution: The maintenance port generates, uses, and can radiate radio-frequency energy through its cables. This energy may interfere with radio and television reception. Disconnect all maintenance port cables when not communicating with the controller through the local connection.

Setting Up a Single Controller


Powering On and Establishing Communication

1. Connect the computer or terminal to the controller, as shown in [Figure 36](#).
The connection to the computer is through the COM1 or COM2 port.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Verify that the computer or terminal is configured as follows:
 - 9600 baud
 - 8 data bits
 - 1 stop bit
 - no parity
 - no flow control
5. Press **Enter**. A copyright notice and the CLI prompt is displayed, indicating that you established a local connection with the controller.

Cabling a Single Controller

The cabling for a single controller is shown in [Figure 37](#).

Note: It is a good idea to plug only the controller cables into the switch. The host cables are plugged into the switch as part of the configuration procedure (["Configuring a Single Controller Using CLI"](#), page 136).



- | | |
|--|--|
| ① Controller
② Host port 1
③ Host port 2 | ④ Cable from the switch to the host Fibre Channel adapter
⑤ FC switch |
|--|--|

Figure 37: Single controller cabling

Configuring a Single Controller Using CLI

To configure a single controller using CLI involves the following processes:

- ["Establishing Controller Connection"](#), page 136
- ["Configuring Controller Settings"](#), page 137
- ["Setting This Controller"](#), page 137
- ["Setting Time and Verifying All Commands"](#), page 140
- ["Plugging in the FC Cable and Verifying Connections"](#), page 143
- ["Repeating Procedure for Each Host Adapter"](#), page 144

Establishing Controller Connection

1. Establish a serial data connection to the RAID system Controller A from your host.
2. From the CLI prompt, enter:

```
SHOW THIS_CONTROLLER FULL
```

The current configuration of your RAID HSG80 controller opens.

Configuring Controller Settings

3. Set the SCSI version using the following command syntax:

```
SET THIS SCSI_VERSION=SCSI-2, or
```

```
SET THIS SCSI_VERSION=SCSI-3
```

Note: If SCSI-2 is selected, you must disable CCL using the command:

```
SET THIS NOCOMMAND_CONSOLE_LUN
```

Note: HP recommends that SCSI-3 mode be used.

4. Check the configuration (parameter settings) of your controller against those listed in [Table 14](#), which are the correct HSG80 controller parameter settings for an IBM-AIX environment.

Table 14: Correct controller parameter settings

HSG80 Parameter Settings for:	Loop	Fabric
Prompt	HSG80>	HSG80>
Node (WWN)	nnnn-nnnn-nnnn-nnnn	nnnn-nnnn-nnnn-nnnn
Command Console LUN	Enabled	Enabled
Port 1 Topology	LOOP HARD	FABRIC
Port 2 Topology	LOOP HARD	FABRIC
Port 1 ALPA	D5	N/A
Port 2 ALPA	D6	N/A
SCSI Version	SCSI-2 or SCSI-3	SCSI-2 or SCSI-3
New Connection OS Type	AIX_CAMBEX	AIX_CAMBEX

Setting This Controller

5. To set the prompt to the HSG80 controller, enter:

```
SET THIS_CONTROLLER PROMPT="HSG>"
```

6. Clear any persistent error messages by entering the following command:

```
HSG> CLEAR CLI
```

7. Set the node by acquiring the World Wide Name (WWN) node number and checksum from the enclosed label (located on the top surface of the controller enclosure), and by entering the following command:

```
HSG> SET THIS_CONTROLLER NODE=NNNN-NNNN-NNNN-NNNN CHECKSUM
```

8. Enable the command console LUN by entering the following command:

```
HSG> SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

9. Set the port topology offline by entering the following command:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=OFFLINE
```

```
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=OFFLINE
```

10. Set the port topology by entering the following command:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=FABRIC
```

```
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

See [Table 15](#) for an example of LOOP_HARD topology setup.

11. Set the SCSI version by entering the following command:

```
HSG> SET THIS_CONTROLLER SCSI VERSION= <SCSI-2 or SCSI-3>
```

12. Determine if the operating system set for the online controller connection is AIX_CAMBEX, by entering the following command:

```
HSG> SHOW CONNECTIONS
```

Note: Although an IBM operating system entry exists in the CLI and SWCC, for the fibre channel connection, use setting **AIX_CAMBEX** for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

```
HSG> SET !NEWCONxx OPERATING_SYSTEM=AIX_CAMBEX
```

where: xx is the online connection ID produced in the `SHOW CONNECTIONS` command above.

Note: Further information on command syntax can be found in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide*.

13. Restart the controllers for the new settings to take effect, by entering the following command:

```
HSG> RESTART OTHER_CONTROLLER  
HSG> RESTART THIS_CONTROLLER
```

14. Change any controller parameter settings inconsistent with those in [Table 14](#) by entering the appropriate command for that parameter, as described in the following steps.

Note: Ignore the Restart Controller messages until you have finished making all the required configuration parameter changes.

Note: After completing all of the required parameter changes, you must restart the controller for the new settings to take effect.

15. To set the prompt to the HSG80 controller, enter:

```
SET THIS_CONTROLLER PROMPT="HSG>"
```

16. Clear the persistent error messages by entering the following command:

```
HSG> CLEAR CLI
```

17. Set the node by acquiring the Worldwide Name (WWN) node number and checksum from the enclosed label (located on the top surface of the controller enclosure), and by entering the following command:

```
HSG> SET THIS_CONTROLLER NODE=NNNN-NNNN-NNNN-NNNN CHECKSUM
```

18. Enable the command console LUN by entering the following command:

```
HSG> SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

19. Set the port topology offline by entering the following command:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=OFFLINE  
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=OFFLINE
```

20. Set the port topology by entering the following command:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=FABRIC  
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

See [Table 14](#) for an example of LOOP_HARD topology setup.

21. Set the SCSI version by entering the following command:

```
HSG> SET THIS_CONTROLLER SCSI VERSION= <SCSI-2 or SCSI-3>
```

22. Determine if the operating system set for the online controller connection is AIX_CAMBEX, by entering the following command:

```
HSG> SHOW CONNECTIONS
```

Although an IBM operating system entry exists in the CLI and SWCC for the fibre channel connection, use default setting AIX_CAMBEX for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover. If so, use:

```
HSG> SET !NEWCONxx OPERATING_SYSTEM=AIX_CAMBEX
```

where: xx is the online connection ID produced in the SHOW CONNECTIONS command above.

Note: Further information on command syntax can be found in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide*.

23. Restart the controllers for the new settings to take effect, by entering the following command:

```
HSG> RESTART OTHER_CONTROLLER
```

```
HSG> RESTART THIS_CONTROLLER
```

It takes about one minute for the CLI prompt to come back after a RESTART command.

Setting Time and Verifying All Commands

1. Set the time on the controller by entering the following syntax:

```
SET THIS TIME=DD-MMM-YYYY:HH:MM:SS
```

2. Use the FRUTIL utility to set up the battery discharge timer. Enter the following command to start FRUTIL:

```
HSG > RUN FRUTIL
```

When FRUTIL asks if you intend to replace the battery, answer **Y**:

```
Do you intend to replace this controller's cache battery? Y/N  
[N] Y
```

FRUTIL prints out a procedure, but does not give you a prompt. Ignore the procedure and press the **Enter** key.

3. Set up any additional optional controller settings, such as changing the CLI prompt. See the *SET THIS CONTROLLER/OTHER CONTROLLER* command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide* for the format of optional settings.
4. Verify that all commands have taken effect. Use the following command:

HSG> SHOW THIS

Verify node ID, allocation class, SCSI version, failover mode, identifier, and port topology.

The following sample is a result of a `SHOW THIS` command, with the areas of interest in bold.

Controller:

HSG80 ZG94214134 Software V8.8, Hardware 0000

NODE_ID = 5000-1FE1-0007-9750

ALLOCATION_CLASS = 0

SCSI_VERSION = SCSI-3

Configured for dual-redundancy with ZG9421461

In dual-redundant configuration

Device Port SCSI address 7

Time: 10-Mar-2002:12:30:34

Command Console LUN is disabled

Smart Error Eject Disabled

Host PORT_1:

Reported PORT_ID = 5000-1FE1-0007-9751

PORT_1_TOPOLOGY = FABRIC (fabric up)

Address = 7D4000

Host PORT_2:

Reported PORT_ID = 5000-1FE1-0007-9752

PORT_2_TOPOLOGY = FABRIC (standby)

Address = 210513

NOREMOTE_COPY

Cache:

512 megabyte write cache, version 0022

Cache is GOOD

No unflushed data in cache

CACHE_FLUSH_TIMER = DEFAULT (10 seconds)

Mirrored Cache:

Not enabled

Battery:

NOUPS

FULLY CHARGED

Expires: 25-JUN-2003

-
5. Turn on the switches, if not done previously.

Plugging in the FC Cable and Verifying Connections

6. Plug the Fibre Channel cable from the first host bus adapter into the switch. Enter the SHOW CONNECTIONS command to view the connection table:

```
HSG> SHOW CONNECTIONS
```

Note: Although an IBM operating system entry exists in the CLI and SWCC for the Fibre Channel connection, use the setting AIX_CAMBEX for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

The first connection has one or more entries in the connection table. Each connection has a default name of the form !NEWCONxx, where xx is a number representing the order in which the connection was added to the connection table.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON00	WINNT	OTHER	1	210513	0L other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

For an explanation about why there may be multiple connections shown for one adapter, see "[Determining Connections](#)", page 41.

7. Rename the connections to something meaningful to the system and easy to remember. For example, to assign the name ANGEL1A1 to connection !NEWCON01, enter:

```
RENAME !NEWCON01 ANGEL1A1
```

For a recommended naming convention, see "[Naming Connections](#)", page 41.

8. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=AIX_CAMBEX
```

Note: Although an IBM operating system entry exists in the CLI and SWCC, for the Fibre Channel connection, use the setting **AIX_CAMBEX** for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

9. Verify the changes:

SHOW CONNECTIONS

Mark or tag all Fibre Channel cables at both ends for ease of maintenance.

Repeating Procedure for Each Host Adapter

10. Repeat step 7, 8, and 9 for each of that adapter's host connections, or delete the unused connections from the table.
11. For each host adapter, repeat steps 6 through 10.

Setting Up a Controller Pair

The following procedures describe how to set up a controller pair.

Powering Up and Establishing Communication

1. Connect the computer or terminal to the controller as shown in [Figure 36](#). The connection to the computer is through the COM1 or COM2 ports.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Configure the computer or terminal as follows:
 - 9600 baud
 - 8 data bits
 - 1 stop bit
 - no parity
 - no flow control
5. Press **Enter**.


A copyright notice and the CLI prompt is displayed, indicating that you established a local connection with the controller.

Cabling a Controller Pair

The cabling for a controller pair is shown in [Figure 38](#).

Note: It is a good idea to plug only the controller cables into the switch or hub. The host cables are plugged into the or hub as part of the configuration procedure ("Configuring a Controller Pair Using CLI", page 146).

Figure 38 shows a controller pair with failover cabling showing one HBA per server with HSG80 controller in transparent failover mode.



CXO6887B

- ① Controller A
- ② Controller B
- ③ Host port 1
- ④ Host port 2
- ⑤ Cable from the switch to the host FC adapter
- ⑥ FC switch

Figure 38: Controller pair failover cabling

Configuring a Controller Pair Using CLI

To configure a controller pair using CLI involves the following processes:

- ["Establishing a Connection to the System Controller"](#), page 146
- ["Configuring Controller Settings"](#), page 147
- ["Setting This Controller"](#), page 148
- ["Setting Time and Verifying All Commands"](#), page 149
- ["Plugging in the FC Cable and Verifying Connections"](#), page 151

Establishing a Connection to the System Controller

1. Establish a serial data connection to RAID system Controller A from your host.
2. From the CLI prompt, enter:

`SHOW THIS_CONTROLLER FULL` The current configuration of your RAID HSG80 controller is displayed.

Configuring Controller Settings

- Check the configuration (parameter settings) of your controller against those listed in [Table 15](#).

Table 15: Correct controller parameter settings, AIX

HSG80 Parameter Settings for:	Loop	Fabric
Prompt	HSG80>	HSG80>
Node (WWN)	nnnn-nnnn-nnnn-nnnn	nnnn-nnnn-nnnn-nnnn
Command Console LUN	Enabled	Enabled
Port 1 Topology	LOOP HARD	FABRIC
Port 2 Topology	LOOP HARD	FABRIC
Port 1 ALPA	D5	N/A
Port 2 ALPA	D6	N/A
SCSI Version	SCSI-2 or SCSI-3	SCSI-2 or SCSI-3
New Connection OS Type	AIX_CAMBEX	AIX_CAMBEX

Note: Setting the SCSI version to SCSI-3 does not make the controller fully compliant with the SCSI-3 standards.

- Change any controller parameter settings inconsistent with those in [Table 15](#) by entering the appropriate command for that parameter, as described in the next section.

Note: After completing all of the required parameter changes, you must restart the controller for the new settings to take effect.

Ignore the Restart Controller messages until you have finished making all the required configuration parameter changes.

Setting This Controller

5. To set the prompt to the HSG80 controller, enter:

```
SET THIS_CONTROLLER PROMPT="HSG>"
```

6. Clear any persistent error messages by entering the following command:

```
HSG> CLEAR CLI
```

7. Set the node by acquiring the World Wide Name (WWN) node number and checksum from the enclosed label (located on the top surface of the controller enclosure), and by entering the following command:

```
HSG> SET THIS_CONTROLLER NODE=NNNN-NNNN-nnnn-NNNN CHECKSUM
```

8. Enable the command console LUN by entering the following command:

```
HSG> SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

9. Set the port topology offline by entering the following command:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=OFFLINE
```

```
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=OFFLINE
```

10. Set the port topology by entering the following command:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=FABRIC
```

```
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

See [Table 15](#) for an example of LOOP_HARD topology setup.

11. Set the SCSI version by entering the following command:

```
HSG> SET THIS_CONTROLLER SCSI VERSION= <SCSI-2 or SCSI-3>
```

12. Determine if the operating system set for the online controller connection is AIX_CAMBEX, by entering the following command:

```
HSG> SHOW CONNECTIONS
```

Note: Although an IBM operating system entry exists in the CLI and SWCC, for the fibre channel connection, use setting **AIX_CAMBEX** for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

```
HSG> SET !NEWCONxx OPERATING_SYSTEM=AIX_CAMBEX
```

where: **xx** is the online connection ID produced in the **SHOW CONNECTIONS** command above.

Note: Further information on command syntax can be found in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide*.

13. Restart the controllers for the new settings to take effect, by entering the following command:

```
HSG> RESTART OTHER_CONTROLLER
```

```
HSG> RESTART THIS_CONTROLLER
```

It takes about a minute for the CLI prompt to come back after a RESTART command.

Setting Time and Verifying All Commands

14. Set the time on the controller by entering the following syntax:

```
SET THIS TIME=DD-MMM-YYYY:HH:MM:SS
```

15. Use the FRUTIL utility to set up the battery discharge timer. Enter the following command to start FRUTIL:

```
HSG> RUN FRUTIL
```

When FRUTIL asks if you intend to replace the battery, answer **Y**:

```
Do you intend to replace this controller's cache battery? Y/N  
[N] Y
```

FRUTIL prints out a procedure, but does not give you a prompt. Ignore the procedure and press **Enter**.

16. Set up any additional optional controller settings, such as changing the CLI prompt. See the SET THIS CONTROLLER/OTHER CONTROLLER command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide* for the format of optional settings. Perform this step on both controllers.
17. Verify that all commands have taken effect by entering the following command:

```
HSG> SHOW THIS
```

18. Verify node ID, allocation class, SCSI version, failover mode, identifier, and port topology. The following display is a sample result of a **SHOW THIS** command, with the areas of interest in bold.

Controller:

```
HSG80 ZG94214134 Software V8.8, Hardware 0000
NODE_ID          = 5000-1FE1-0007-9750
ALLOCATION_CLASS = 0
SCSI_VERSION       = SCSI-3
Configured for dual-redundancy with ZG9421461
In dual-redundant configuration
Device Port SCSI address 7
Time: 10-Mar-2002:12:30:34
Command Console LUN is disabled
Smart Error Eject Disabled
```

Host PORT_1:

```
Reported PORT_ID = 5000-1FE1-0007-9751
PORT_1_TOPOLOGY = FABRIC (fabric up)
Address          = 7D4000
```

Host PORT_2:

```
Reported PORT_ID = 5000-1FE1-0007-9752
PORT_2_TOPOLOGY = FABRIC (standby)
Address          = 210513
```

NOREMOTE_COPY

Cache:

```
512 megabyte write cache, version 0022
Cache is GOOD
No unflushed data in cache
CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
```

Mirrored Cache:

Not enabled

Battery:

NOUPS

FULLY CHARGED

Expires: 25-JUN-2003

19. Turn on the switches, if not done previously.

Plugging in the FC Cable and Verifying Connections

20. Plug the FC cable from the first host adapter into the switch or hub. Enter a SHOW CONNECTIONS command to view the connection table:

```
HSG > SHOW CONNECTIONS
```

The first connection has one or more entries in the connection table. Each connection has a default name of the form !NEWCONxx, where xx is a number representing the order in which the connection was added to the connection table.

For a description of why plugging in one adapter can result in multiple connections, see "[Numbers of Connections](#)", page 42.

21. Rename the connections to something meaningful to the system and easy to remember. For example, to assign the name ANGEL1A1 to connection !NEWCON01, enter:

```
HSG > RENAME !NEWCON01 ANGEL1A1
```

HP recommends using a naming convention, see "[Naming Connections](#)", page 41.

22. Specify the operating system for the connection:

```
HSG > SET ANGEL1A1 OPERATING_SYSTEM=AIX_CAMBEX
```

Note: Although an IBM operating system entry exists in the CLI and SWCC, for the Fibre Channel connection, use the default setting **AIX_CAMBEX** for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

23. Verify the changes:

```
HSG > SHOW CONNECTIONS
```

Mark or tag all Fibre Channel cables at both ends for ease of maintenance.

Repeating Procedure for Each Host Adapter Connection

24. Repeat steps 21, 22, and 23 for each of that adapter's host connections or delete the unwanted connections from the table.
25. For each host adapter, repeat steps 20 through 24.

Configuring Devices

The disks on the device bus of the HSG80 can be configured manually or with the CONFIG utility. The CONFIG utility is easier. Invoke CONFIG with the following command:

RUN CONFIG



WARNING: HP recommends that you use the CONFIG utility only at reduced I/O loads.

The CONFIG utility takes about two minutes to discover and to map the configuration of a completely populated storage system.

Configuring Storage Containers

For a technology refresher on this subject, refer to "[Choosing a Container Type](#)", page 76.

In choosing a container, you choose between independent disks (JBODs) or one of several storageset types, as shown in [Figure 39](#). The independent disks and the selected storageset may also be partitioned.

The following configurations are detailed in the following sections:

- ["Configuring a Stripeset"](#), page 154
- ["Configuring a Mirrorset"](#), page 154
- ["Configuring a RAIDset"](#), page 155
- ["Configuring a Striped Mirrorset"](#), page 156
- ["Configuring a Single-Disk Unit \(JBOD\)"](#), page 157
- ["Configuring a Partition"](#), page 157

Containers




Figure 39: Storage container types

Configuring a Stripeset

1. Create the stripeset by adding its name to the controller's list of storage sets and by specifying the disk drives it contains. Use the following syntax:

```
ADD STRIPESET STRIPESET-NAME DISKNNNNNN DISKNNNNNN . . . . .
```

2. Initialize the stripeset, specifying any desired switches:

```
INITIALIZE STRIPESET-NAME SWITCHES
```

See "[Specifying Initialization Switches](#)", page 93, for a description of the initialization switches.

3. Verify the stripeset configuration:

```
SHOW STRIPESET-NAME
```

4. Assign the stripeset a unit number to make it accessible by the hosts. See "[Assigning Unit Numbers and Unit Qualifiers](#)", page 159.

For example:

The commands to create Stripe1, a stripeset consisting of three disks (DISK10000, DISK20000, and DISK10100) and having a chunksize of 128:

```
ADD STRIPESET STRIPE1 DISK10000 DISK20000 DISK30000
```

```
INITIALIZE STRIPE1 CHUNKSIZE=128
```

```
SHOW STRIPE1
```

Configuring a Mirrorset

1. Create the mirrorset by adding its name to the controller's list of storage sets and by specifying the disk drives it contains. Optionally, you can append mirrorset switch values:

```
ADD MIRRORSET MIRRORSET-NAME DISKNNNNNN DISKNNNNNN SWITCHES
```

Note: See the ADD MIRRORSET command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide* for a description of the mirrorset switches.

2. Initialize the mirrorset, specifying any desired switches:

```
INITIALIZE MIRRORSET-NAME SWITCHES
```

See "[Specifying Initialization Switches](#)", page 93, for a description of the initialization switches.

3. Verify the mirrorset configuration:

```
SHOW MIRRORSET-NAME
```

4. Assign the mirrorset a unit number to make it accessible by the hosts. See "[Assigning Unit Numbers and Unit Qualifiers](#)", page 159.

For example:

The commands to create Mirr1, a mirrorset with two members (DISK10000 and DISK20000), and to initialize it using default switch settings:

```
ADD MIRRORSET MIRR1 DISK10000 DISK20000  
INITIALIZE MIRR1  
SHOW MIRR1
```

Configuring a RAIDset

1. Create the RAIDset by adding its name to the controller's list of storagesets and by specifying the disk drives it contains. Optionally, you can specify RAIDset switch values:

```
ADD RAIDSET RAIDSET-NAME DISKNNNNN DISKNNNNN DISKNNNNN SWITCHES
```

Note: See the `ADD RAIDSET` command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide* for a description of the RAIDset switches.

2. Initialize the RAIDset, specifying any desired switches:

```
INITIALIZE RAIDSET-NAME SWITCH
```

Note: HP recommends that you allow initial reconstruct to complete before allowing I/O to the RAIDset. Not doing so may generate forced errors at the host level. To determine whether initial reconstruct has completed, enter `SHOW RAIDSET FULL`.

See "[Specifying Initialization Switches](#)", page 93, for a description of the initialization switches.

3. Verify the RAIDset configuration:

```
SHOW RAIDSET-NAME
```

4. Assign the RAIDset a unit number to make it accessible by the hosts. See "[Assigning Unit Numbers and Unit Qualifiers](#)", page 159.

For example:

The commands to create RAID1, a RAIDset with three members (DISK10000, DISK20000, and DISK10100) and to initialize it with default values:

```
ADD RAIDSET RAID1 DISK10000 DISK20000 DISK30000  
INITIALIZE RAID1  
SHOW RAID1
```

Configuring a Striped Mirrorset

1. Create, but do not initialize, at least two mirrorsets.

See "[Configuring a Mirrorset](#)", page 154.

2. Create a stripeset and specify the mirrorsets it contains:

```
ADD STRIPESET STRIPESET-NAME MIRRORSET-1  
MIRRORSET-2 . . . MIRRORSET-N
```

3. Initialize the striped mirrorset, specifying any desired switches:

```
INITIALIZE STRIPESET-NAME SWITCH
```

See "[Specifying Initialization Switches](#)", page 93, for a description of the initialization switches.

4. Verify the striped mirrorset configuration:

```
SHOW STRIPESET-NAME
```

5. Assign the stripeset mirrorset a unit number to make it accessible by the hosts. See "[Assigning Unit Numbers and Unit Qualifiers](#)", page 159.

For example:

The commands to create Stripe1, a striped mirrorset that comprises Mirr1, Mirr2, and Mirr3, each of which is a two-member mirrorset:

```
ADD MIRRORSET MIRR1 DISK10000 DISK20000  
ADD MIRRORSET MIRR2 DISK20100 DISK10100  
ADD MIRRORSET MIRR3 DISK10200 DISK20200  
ADD STRIPESET STRIPE1 MIRR1 MIRR2 MIRR3  
INITIALIZE STRIPE1  
SHOW STRIPE1
```

Configuring a Single-Disk Unit (JBOD)

1. Initialize the disk drive, specifying any desired switches:

`INITIALIZE DISK-NAME SWITCHES`

See "[Specifying Initialization Switches](#)", page 93, for a description of the initialization switches.

2. Verify the configuration by entering the following command:

`SHOW DISK-NAME`

3. Assign the disk a unit number to make it accessible by the hosts. See "[Assigning Unit Numbers and Unit Qualifiers](#)", page 159.

Configuring a Partition

1. Initialize the storageset or disk drive, specifying any desired switches:

`INITIALIZE STORAGESET-NAME SWITCHES`

or

`INITIALIZE DISK-NAME SWITCHES`

See "[Specifying Initialization Switches](#)", page 93, for a description of the initialization switches.

2. Create each partition in the storageset or disk drive by indicating the partition's size. Also specify any desired switch settings:

`CREATE_PARTITION STORAGESET-NAME SIZE=N SWITCHES`

or

`CREATE_PARTITION DISK-NAME SIZE=N SWITCHES`

where N is the percentage of the disk drive or storageset that is assigned to the partition. Enter `SIZE=LARGEST`, on the last partition only, to let the controller assign the largest free space available to the partition.

Note: See the `CREATE_PARTITION` command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide* for a description of the partition switches.

3. Verify the partitions:

`SHOW STORAGESET-NAME`

or

SHOW *DISK-NAME*

The partition number is displayed in the first column, followed by the size and starting block of each partition.

4. Assign the partition a unit number to make it accessible by the hosts. See "[Assigning Unit Numbers and Unit Qualifiers](#)", page 159.

For example:

The commands to create RAID1, a three-member RAIDset, then partition it into two storage units are shown below.

```
ADD RAIDSET RAID1 DISK10000 DISK20000 DISK30000
INITIALIZE RAID1
CREATE_PARTITION RAID1 SIZE=25
CREATE_PARTITION RAID1 SIZE=LARGEST
SHOW RAID1
```

Assigning Unit Numbers and Unit Qualifiers

Each storageset, partition, or single (JBOD) disk must be assigned a unit number for the host to access. As the units are added, their properties can be specified through the use of command qualifiers, which are discussed in detail under the ADD UNIT command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface Reference Guide*.

Because of different SCSI versions, refer to the section "["Assigning Unit Numbers Depending on SCSI_VERSION"](#)", page 49. The choice for SCSI_VERSION effects how certain unit numbers and host connection offsets interact.

Each unit can be reserved for the exclusive use of a host or group of hosts. See "["Restricting Host Access in Transparent Failover Mode"](#)", page 51, and "["Restricting Host Access in Multiple-Bus Failover Mode"](#)", page 55.

Assigning a Unit Number to a StorageSet

To assign a unit number to a storageset, use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESSET-NAME
```

For example:

To assign unit D102 to RAIDset R1, use the following command:

```
ADD UNIT D102 R1
```

Assigning a Unit Number to a Single (JBOD) Disk

To assign a unit number to a single (JBOD) disk, use the following syntax:

```
ADD UNIT UNIT-NUMBER DISK-NAME
```

For example:

To assign unit D4 to DISK20300, use the following command:

```
ADD UNIT D4 DISK20300
```

Assigning a Unit Number to a Partition

To assign a unit number to a partition, use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESSET-NAME PARTITION=PARTITION-NUMBER
```

For example:

To assign unit D100 to partition 3 of mirrorset mirr1, use the following command:

```
ADD UNIT D100 MIRR1 PARTITION=3
```

Configuration Options

There are multiple options that allow you to configure your system.

Changing the CLI Prompt

To change the CLI prompt, enter a 1- to 16- character string as the new prompt, according to the following syntax:

```
SET THIS_CONTROLLER PROMPT = "NEW PROMPT"
```

If you are configuring dual-redundant controllers, also change the CLI prompt on the “other controller.” Use the following syntax:

```
SET OTHER_CONTROLLER PROMPT = "NEW PROMPT"
```

Note: It is suggested that the prompt name reflect some information about the controllers. For example, if the subsystem is the third one in a lab, name the top controller prompt, LAB3A and the bottom controller, LAB3B.

Mirroring Cache

To specify mirrored cache, use the following syntax:

```
SET THIS MIRRORED_CACHE
```

Adding Disk Drives

If you add new disk drives to the subsystem, the disk drives must be added to the controllers’ list of known devices:

- To add one new disk drive to the list of known devices, use the following syntax:

```
ADD DISK DISKNNNNN P T L
```

- To add several new disk drives to the list of known devices, enter the following command:

```
RUN CONFIG
```

Adding a Disk Drive to the Spareset

The spareset is a collection of spare disk drives that are available to the controller should it need to replace a failed member of a RAIDset or mirrorset.

Note: This procedure assumes that the disks that you are adding to the spareset have already been added to the controller's list of known devices.

To add the disk drive to the controller's spareset list, use the following syntax:

```
ADD SPARESET DISKNNNNN
```

Repeat this step for each disk drive you want to add to the spareset:

For example:

The following example shows the syntax for adding DISK11300 and DISK21300 to the spareset.

```
ADD SPARESET DISK11300
```

```
ADD SPARESET DISK21300
```

Removing a Disk Drive from the Spareset

You can delete disks in the spareset if you need to use them elsewhere in your subsystem.

1. Show the contents of the spareset entering the following command:

```
SHOW SPARESET
```

2. Delete the desired disk drive entering the following command:

```
DELETE SPARESET DISKNNNNN
```

The RUN CONFIG command does not delete disks from the controllers' device table if a disk has been physically removed or replaced. In this case, you must use the command: DELETE DISKNNNNN.

3. Verify the contents of the spareset by entering the following command:

```
SHOW SPARESET
```

Enabling Autospare

With AUTOSPARE enabled on the failedset, any new disk drive that is inserted into the PTL location of a failed disk drive is automatically initialized and placed into the spareset. If initialization fails, the disk drive remains in the failedset until you manually delete it from the failedset.

To enable autospare, use the following command:

```
SET FAILEDSET AUTOSPARE
```

To disable autospare, use the following command:

```
SET FAILEDSET NOAUTOSPARE
```

During initialization, AUTOSPARE checks to see if the new disk drive contains metadata. Metadata is information the controller writes on the disk drive when the disk drive is configured into a storageset. Therefore, the presence of metadata indicates that the disk drive belongs to, or has been used by, a storageset. If the disk drive contains metadata, initialization stops. (A new disk drive does not contain metadata but a repaired or reused disk drive might. To erase metadata from a disk drive, add it to the controller's list of devices, then set it to be nontransportable and initialize it.)

Deleting a Storageset

Note: If the storageset you are deleting is partitioned, you must delete each partitioned unit before you can delete the storageset.

1. Show the storageset's configuration:

```
SHOW STORAGESET-NAME
```

2. Delete the unit number that uses the storageset. Use the following command:

```
DELETE UNIT-NUMBER
```

3. Delete the storageset. Use the following command:

```
DELETE STORAGESET-NAME
```

4. Verify the configuration:

```
SHOW STORAGESET-NAME
```

Changing Switches for a Storageset or Device

You can optimize a storageset or device at any time by changing the switches that are associated with it. Remember to update the storageset profile when changing its switches.

Displaying the Current Switches

To display the current switches for a storageset or single-disk unit, enter a SHOW command, specifying the FULL switch:

```
SHOW STORAGESET-NAME
```

or

```
SHOW DEVICE-NAME
```

Note: FULL is not required when showing a particular device. It is used when showing all devices, for example, SHOW DEVICES FULL.

Changing RAIDset and Mirrorset Switches

Use the SET *storageset-name* command to change the RAIDset and Mirrorset switches associated with an existing storageset.

For example, the following command changes the replacement policy for RAIDset RAID1 to BEST_FIT:

```
SET RAID1 POLICY=BEST_FIT
```

Changing Device Switches

Use the SET *device-name* command to change the device switches.

For example, to request a data transfer rate of 20 MHz for DISK10000:

```
SET DISK10000 TRANSFER_RATE_REQUESTED=20MHZ
```

Changing Initialize Switches

The initialization switches cannot be changed without destroying the data on the storageset or device. These switches are integral to the formatting and can only be changed by reinitializing the storageset. Initializing a storageset is similar to formatting a disk drive; all data is destroyed during this procedure.

Changing Unit Switches

Use the SET *unit-name* command to change the characteristics of a unit.

For example, the following command enables write protection for unit D100:

```
SET D100 WRITE_PROTECT
```

Verifying Storage Configuration from Host

Once the HP StorageWorks Solution Software has been installed, logical units may be created on the RAID Array and configured. Create these logical units using the Client or the CLI. Be sure the Agent and Client are already installed on the AIX and Windows systems, respectively.

New logical units created on the RAID Array must be recognized by the operating system before they can be used. The following steps outline the process:

1. Using the CLI or the Client, create the logical units on the RAID Array.
2. As the superuser (*root*), run the command on the host:

```
# cfgmgr -v
```

3. Verify that the host has located and configured the new RAID Array devices in the operating system using the command:

```
# lshsg
```

The output of the command looks like:

```
hdisk1 Available X0-06-01-0,0 DEC HSG80 Command Console LUN
hdisk2 Available X0-06-01-0,1 DEC HSG80 RAID Array
hdisk3 Available X0-06-01-0,2 DEC HSG80 RAID Array
hdisk4 Available X0-06-01-0,3 DEC HSG80 RAID Array
hdisk5 Available X0-06-01-0,4 DEC HSG80 RAID Array
hdisk6 Available X0-06-01-0,5 DEC HSG80 RAID Array
hdisk7 Available X0-06-01-0,6 DEC HSG80 RAID Array
hdisk8 Available X0-06-01-0,7 DEC HSG80 RAID Array
hdisk9 Available X0-06-01-2,2 DEC HSG80 RAID Array
hdisk10 Available X0-06-01-2,4 DEC HSG80 RAID Array
hdisk11 Available X0-06-01-4,4 DEC HSG80 RAID Array
hdisk12 Available X0-06-01-5,1 DEC HSG80 RAID Array
```

You should have an entry of type:

```
hdisknn Available X0-06-01-2,4 DEC HSG80 RAID Array
```

for each logical unit that you have created on the RAID Array.

4. If the devices were not recognized, recheck the status of the RAID controller, logical units, and cable connections.

You can now use the **System Management Interface Tool (SMIT)**, to make volume groups, logical volumes and file systems using these logical units.

Using CLI for Configuration

6

This chapter presents an example of how to configure a storage subsystem using the Command Line Interpreter (CLI). The CLI configuration example shown assumes:

- A normal, new controller pair, which includes:
 - NODE ID set
 - No previous failover mode
 - No previous topology set
- Two single-bus model 4214R disk enclosure shelves
- PCMCIA cards installed in both controllers


A storage subsystem example is shown in [Figure 40](#). The example system contains three non-clustered AIX hosts, as shown in [Figure 41](#). The resulting virtual system, from the host's point of view, is shown in [Figure 42](#).

[Figure 40](#) shows an example storage system map for the BA370 enclosure. Details on building your own map are described in [Chapter 2](#). Templates to help you build your storage map are supplied in [Appendix A](#).

Port								Targets
	1	2	3	4	5	6		
Power Supply	D2 S2 DISK103 00	D2 S2 DISK203 00	D2 S2 DISK303 00	D2 S2 DISK403 00	D2 D101 DISK503 00	spareset member DISK603 00	Power Supply	3
Power Supply	D0 S1 MI DISK102 00	D0 S1 M1 DISK202 00	D0 S1 M2 DISK302 00	D0 S1 M2 DISK402 00	D1 M3 DISK502 00	D1 M3 DISK602 00	Power Supply	2
Power Supply	D120 R2 DISK101 00	D120 R2 DISK201 00	D120 R2 DISK301 00	D120 R2 DISK401 00	D120 R2 DISK501 00	D120 R2 DISK601 00	Power Supply	1
Power Supply	D102 R1 DISK100 00	D102 R1 DISK200 00	D102 R1 DISK300 00	D102 R1 DISK400 00	D102 R1 DISK500 00	D102 R1 DISK600 00	Power Supply	0

Figure 40: Example storage map for the BA370 enclosure

The example system, shown in [Figure 41](#), contains three non-clustered AIX hosts. Port 1 link is separate from port 2 link (that is, ports 1 of both controllers are on one loop or fabric, and port 2 of both controllers are on another). Therefore, each adapter has two connections. The target numbers represent what the third number of the disk number is. For example, if the disk number is disk10200, 2 is the target number.




NOTE: FCA = Fibre Channel Adapter

CXO7547B

Figure 41: Example, three non-clustered host systems


[Figure 42](#) represents units that are logical or virtual disks comprised of storage sets configured from physical disks.

[Figure 43](#) represents an example of virtual system layout from the hosts' point of view.



CXO7110B

Figure 42: Example, logical or virtual disks comprised of storage sets



CXO7297B

Figure 43: Example, virtual system layout from hosts' point of view

CLI Configuration Example

Text conventions used in this example are listed below:

- Text in *italics* indicates an action you take.
- Text in **THIS FORMAT**, indicates a command you type. Be certain to press **Enter** after each command.
- Text enclosed within a box, indicates information that is displayed by the CLI interpreter.

Note: "This" controller is top controller (A).

Plug serial cable from maintenance terminal into top controller.

```
CLEAR CLI  
SET FAILOVER COPY=THIS  
CLEAR CLI  
SET THIS SCSI_VERSION=SCSI-2  
SET THIS ALLOCATION_CLASS=0  
RESTART OTHER  
RESTART THIS  
SET THIS TIME=10-Mar-2001:12:30:34  
RUN FRUTIL
```

```
Do you intend to replace this controller's cache  
battery? Y/N [Y]
```

Y

Plug serial cable from maintenance terminal into bottom controller.

Note: Bottom controller (B) becomes “this” controller.

```
RUN FRUTIL
```

```
Do you intend to replace this controller's cache  
battery? Y/N [Y]
```

Y

```
SET THIS MIRRORED_CACHE
```

Note: This command causes the controllers to restart.

```
SET THIS PROMPT="HLNDR BOTTOM"  
SET OTHER PROMPT="HLNDR TOP"  
SHOW THIS  
SHOW OTHER
```

Plug in the Fibre Channel cable from the adapter in host “PURPLE.”

```
SHOW CONNECTIONS
```

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON00	WINNT	THIS	1	XXXX XX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XX XX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		

```
RENAME !NEWCON00 PURPLE1A1
SET PURPLE1A1 OPERATING_SYSTEM=AIX_CAMBEX
SHOW CONNECTIONS
```

Note: Connection table sorts alphabetically.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
RED1A1	AIX_CAMBEX	OTHER	1	XXXX XX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XX XX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		

Mark or tag both ends of Fibre Channel cables.

Plug in the Fibre Channel cable from the second adapter in host “WHITE.”

```
SHOW CONNECTIONS
```

Note: The offset for !NEWCON01 is automatically set to 100 because it is plugged into host port 2.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON0 1	WINNT	THIS	2	XXXXXX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		
PURPLE1A 1	AIX_CAMBEX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		

```
RENAME !NEWCON01 WHITE1B2
SET WHITE1B2 OPERATING_SYSTEM=AIX_CAMBEX
SHOW CONNECTIONS
```

Mark or tag both end of Fibre Channel cables.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
PURPLE1A 1	AIX_CAMBEX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		
WHITE1B2	AIX_CAMBEX	THIS	2	XXXXXX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		

Plug in the Fibre Channel cable from the adapter in host “TAN”.

SHOW CONNECTIONS

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON0 2	WINNT	THIS	2	XXXX XX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		
PURPLE1A 1	AIX_CAMBEX	OTHER	1	XXXX XX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		
WHITE1B2	AIX_CAMBEX	THIS	2	XXXX XX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXX X-XXXX		

```
RENAME !NEWCON02 TAN1B2
SET TAN1B2 OPERATING_SYSTEM=AIX_CAMBEX
SET TAN1B2 UNIT_OFFSET=120
SHOW CONNECTIONS
```

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
PURPLE1A 1	AIX_CAMBEX	OTHER	1	XXXX XX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXX X-XXXX				
WHITE1B2	AIX_CAMBEX	THIS	2	XXXX XX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXX X-XXXX				
TAN1B2	AIX_CAMBEX	THIS	2	XXXX XX	OL this	120
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXX X-XXXX				

Mark or tag both end of Fibre Channel cables.

```
RUN CONFIG
ADD RAIDSET R1 DISK10000 DISK20000 DISK30000 DISK40000 DISK50000
DISK60000
INITIALIZE R1
ADD UNIT D102 R1 DISABLE_ACCESS_PATH=ALL
SET D102 ENABLE_ACCESS_PATH=PURPLE1A1
ADD RAIDSET R2 DISK10100 DISK20100 DISK30100 DISK40100 DISK50100
DISK60100
INITIALIZE R2
ADD UNIT D120 R2 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=(TAN1B2)
ADD MIRRORSET MI DISK10200 DISK20200
ADD MIRRORSET M2 DISK30200 DISK40200
ADD STRIPESET S1 M1 M2
INITIALIZE S1
ADD UNIT D0 S1 DISABLE_ACCESS_PATH=ALL
SET D0 ENABLE_ACCESS_PATH=PURPLE1A1
ADD MIRRORSET M3 DISK50200 DISK60200
INITIALIZE M3
ADD UNIT D1 M3 DISABLE_ACCESS_PATH=ALL
SET D1 ENABLE_ACCESS_PATH=(WHITE1B2)
ADD STRIPESET S2 DISK10300 DISK20300 DISK30300 DISK40300
INITIALIZE S2
ADD UNIT D2 S2 DISABLE_ACCESS_PATH=ALL
SET D2 ENABLE_ACCESS_PATH=PURPLE1A1
INITIALIZE DISK50300
ADD UNIT D101 DISK50300 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=(WHITE1B2)
ADD SPARESET DISK60300
SHOW UNITS FULL
```


7

Backing Up, Cloning, and Moving Data

This chapter includes the following topics:

- "Backing Up Subsystem Configurations", page 178
- "Creating Clones for Backup", page 179
- "Clone and Snapshot and AIX Environment", page 183
- "Moving Storagesets", page 191

Backing Up Subsystem Configurations

The controller stores information about the subsystem configuration in its nonvolatile memory. This information could be lost if the controller fails or when you replace a module in the subsystem.

Use the following command to produce a display that shows if the save configuration feature is active and which devices are being used to store the configuration.

```
SHOW THIS_CONTROLLER FULL
```

The resulting display includes a line that indicates status and how many devices have copies of the configuration. The last line shows on how many devices the configuration is backed up.

Note: DO NOT use `SAVE_CONFIGURATION` in dual redundant controller installations. It is not supported and may result in unexpected controller behavior.

The `SHOW DEVICES FULL` command shows which disk drives are set up to back up the configuration. The syntax for this command is shown below:

```
SHOW DEVICES FULL
```

Creating Clones for Backup

Use the Clone utility to duplicate the data on any unpartitioned single-disk unit, stripeset, mirrorset, or striped mirrorset in preparation for backup. When the cloning operation is complete, you can back up the Clones rather than the storageset or single-disk unit, which can continue to service its I/O load. When you are cloning a mirrorset, Clone does not need to create a temporary mirrorset. Instead, it adds a temporary member to the mirrorset and copies the data onto this new member.

The Clone utility creates a temporary, two-member mirrorset for each member in a single-disk unit or stripeset. Each temporary mirrorset contains one disk drive from the unit you are cloning and one disk drive onto which Clone copies the data. During the copy operation, the unit remains online and active so that the Clones contain the most up-to-date data.

After the Clone utility copies the data from the members to the Clones, it restores the unit to its original configuration and creates a Clone unit you can back up. The Clone utility uses the steps shown in [Figure 44](#) to duplicate each member of a unit.




Figure 44: Clone utility steps for duplicating unit members

To Clone a single-disk unit, stripeset, or mirrorset:

1. Establish a connection to the controller that accesses the unit you want to Clone.

2. Start Clone using the following command:

RUN CLONE

3. When prompted, enter the unit number of the unit you want to Clone.
4. When prompted, enter a unit number for the Clone unit that Clone creates.
5. When prompted, indicate how you would like the Clone unit to be brought online: either automatically or only after your approval.
6. When prompted, enter the disk drives you want to use for the Clone units.
7. Back up the Clone unit.

The following example shows the commands you would use to Clone storage unit D6. The Clone command terminates after it creates storage unit D33, a Clone or copy of D6.

```
RUN CLONE  
CLONE LOCAL PROGRAM INVOKED  
UNITS AVAILABLE FOR CLONING:  
98  
ENTER UNIT TO CLONE? 98  
CLONE WILL CREATE A NEW UNIT WHICH IS A COPY OF UNIT 98.  
ENTER THE UNIT NUMBER WHICH YOU WANT ASSIGNED TO THE NEW UNIT?  
99  
THE NEW UNIT MAY BE ADDED USING ONE OF THE FOLLOWING METHODS:  
1. CLONE WILL PAUSE AFTER ALL MEMBERS HAVE BEEN COPIED. THE USER  
MUST THEN PRESS RETURN TO CAUSE THE NEW UNIT TO BE ADDED.  
2. AFTER ALL MEMBERS HAVE BEEN COPIED, THE UNIT WILL BE ADDED  
AUTOMATICALLY.  
UNDER WHICH ABOVE METHOD SHOULD THE NEW UNIT BE ADDED[ ]?1  
DEVICES AVAILABLE FOR CLONE TARGETS:  
DISK20200 (SIZE=832317)  
DISK20300 (SIZE=832317)  
USE AVAILABLE DEVICE DISK20200(SIZE=832317) FOR MEMBER  
DISK10300(SIZE=832317) (Y,N) [Y]? Y  
MIRROR DISK10300 C_MA  
SET C_MA NOPOLICY  
SET C_MA MEMBERS=2  
SET C_MA REPLACE=DISK20200  
DEVICES AVAILABLE FOR CLONE TARGETS:  
DISK20300 (SIZE=832317)  
USE AVAILABLE DEVICE DISK20300(SIZE=832317) FOR MEMBER  
DISK10000(SIZE=832317) (Y,N) [Y]? Y  
MIRROR DISK10000 C_MB  
SET C_MB NOPOLICY  
SET C_MB MEMBERS=2  
SET C_MB REPLACE=DISK20300  
COPY IN PROGRESS FOR EACH NEW MEMBER. PLEASE BE PATIENT...  
.
```

```
COPY FROM DISK10300 TO DISK20200 IS 100% COMPLETE
COPY FROM DISK10000 TO DISK20300 IS 100% COMPLETE
PRESS RETURN WHEN YOU WANT THE NEW UNIT TO BE CREATED
REDUCE DISK20200 DISK20300
UNMIRROR DISK10300
UNMIRROR DISK10000
ADD MIRRORSET C_MA          DISK20200
ADD MIRRORSET C_MB          DISK20300
ADD STRIPESET C_ST1 C_MA C_MB
INIT C_ST1      NODESTROY
ADD UNIT D99 C_ST1
D99 HAS BEEN CREATED. IT IS A CLONE OF D98.
CLONE - NORMAL TERMINATION
```

Clone and Snapshot and AIX Environment

The Clone and Snapshot functions permit an HP StorageWorks storage system to make copies of its internal drives. Once these copies become available, applications such as backup can be implemented on the copies, and not affect production work with the original drives.

In AIX, volume groups can be copied in this manner, however AIX does not permit the existence of two or more exact volume group copies in same system environment. The following describes the steps necessary to make the appropriate copies (Clone or Snapshot) of AIX volume groups and to import these copies into the same environment as the originals.

The basic system requirements are an AIX RS/6000 system configured with an HSG80 storage system with enough unassigned drives to make the copy.

Volume group copies can be imported through the use of the AIX command `recreatevg`. Additionally, in order to use the Snapshot function, the HSG80 controller must be loaded with “S” type firmware.

Note: In the case of more than one candidate volume groups on the same system targeted for copy, this procedure should be executed for each candidate one at a time. The volume group can remain imported and their respective logical volumes remain mounted during a Clone or Snapshot, however, writes to the volume group can break the data consistency while a Clone operation is active. When cloning is complete, then volume group access can resume with the original without affecting the copy. This is not the case for Snapshot, however, and as such you can continue volume group access while also accessing the Snapshot volume group.

In the solution software, there are tools to help create, list, delete Snapshots out on the HSG80 from the operating system. There are also tools in this kit to automatically redefine a volume group of a Snapshot or Clone so that it does not appear to the operating system to be the same volume group as the original.

There are two sections that follow. The first describes the Snapshot tools that are included with the solution software. The second is a description of how to manually create a Snapshot and redefine the volume group for use with AIX.

Snapshot Tools from the AIX Solution Kit

The Snapshot tools from the AIX Solution Kit are contained within a fileset called `hsg_snapshot.rte`. The Snapshot utilities can be accessed through SMIT through the menu items **Applications->Compaq HSG80 Snapshots**. Following is a description of each of the menu items under StorageWorks HSG80 Snapshots and how to use them.

List Snapshot Units

The menu item lists all the Snapshot units on a particular HSG80. You have to specify which communications device you wish to use when you list the Snapshot units. This allows you to specify which subsystem you wish to list. The World Wide Node Name of the subsystem attached to your server is listed along with the communications device.

Automatic Snapshot

This automatically creates a Snapshot of a particular HSG80 hdisk device on your system. The only parameter that you need to specify is which hdisk you wish to create a Snapshot of. The utility automatically creates either a JBOD disk or a Stripeset disk to house the Snapshot data and then create the actually Snapshot itself. You can then issue a `cfgmgr -v` command to configure the new device.

If the Snapshot fails to be created, the following could be the cause:

- There is not enough space left on the HSG80.
- The Stripeset could not be created.
- You have reached the maximum allowed number of Snapshots.
- There was a communications failure with the HSG80.

Manual Snapshot

This method of Snapshot creation allows you to choose the parameters you wish to use to create the Snapshot. You first pick the communications device to use the same way as you do when you list the Snapshot units on a particular subsystem. You then fill in the following parameters:

1. **Unit to Snapshot**—The unit name on the subsystem that you wish to make a Snapshot of.

2. **Name of New Unit**—The unit name (Dxx) that you wish to use for the Snapshot unit. Note that the unit must exist on the same controller as the original virtual disk. So, for example, if you have a unit called d5 you wish to Snapshot, it must be within the range of d1–d99 to be on the same controller.
3. **Disk Devices for Snapshot**—The list of physical disks on the subsystem you wish to use to create the Stripeset for the Snapshot.

The utility takes your parameters and tries to create a Snapshot. If the Snapshot fails it could be either because your parameters were incorrect or for the same reasons that an automatic Snapshot would fail.

Delete Snapshot

This allows you to delete a Snapshot unit out on the subsystem. You choose a communications device the same way as you do to list the Snapshots units or create a manual Snapshot and then choose the unit number of the Snapshot that you wish to delete. The utility then deletes the unit and the underlying virtual disk.

Recreate Volume Group of a Snapshot

The function of this menu item is to redefine volume group of a Snapshot once you have configured the hdisk into the operating system. This menu item actually contains two submenu items:

Show Disks with the Same Physical Volume ID

This menu item displays each HSG80 disk and any disks that have the same physical volume identifier.

Recreate Volume Group

This allows you to recreate the volume group of disk that has the same physical volume identifier as another. The only field to fill in is the name of the disk (from the list of disks presented). The utility then changes the physical volume i.d. on the disk, issue the `recreatevg` command on the disk, and then mount any filesystems on the disk. The new volume group name is in the form `snapX_old volume group name` and any of the filesystems is mounted in the directory `/snapX`. So, for example, if the original volume group is named `testdisk1` and it is mounted over `/testdisk1`, the recreated volume group is named `snap1_testdisk1` and the filesystem is mounted on `/snap1/testdisk1`.

Manual Snapshot/Clone Creation

To Clone or Snapshot drives within an AIX volume group, do the following:

1. Locate all physical volumes (hdisk) within the target volume group using the following command:

```
lsvg -p <Target Volume Group>
```

The response lists the hdisk associated with the target volume group.

2. Get the location codes for the hdisk found with `lspv`.

Refer to the AIX documentation on device locations codes.

Each RAID set discovered by AIX is assigned an hdisk designation along with a location code. The same result occurs for the Command Console LUN (CCL), which must be enabled prior to AIX device discovery. The location code for the `hdisk`s can be found by issuing the following AIX command:

```
lsdev -Cc disk
```

Successful command execution shows a list of hdisk that include the ones reported by the `lsvg` command in Step 1, and which also must be described as “DEC HSG80 RAID ARRAY”. An hdisk described as “DEC HSG80 Command Console LUN” must also display in the list as well which is the CCL.

3. Identify the RAID set units associated with the hdisk reported in Step 1.

For this discussion, a single host adapter connection to the storage system is assumed; a dual host adapter connection description follows.

AIX device location codes for hdisk have the format, AA-BB-CC-S,L, where AA-BB-CC resolves to the AIX system host adapter the storage system is attached to. S is the storage systems port number and is mapped to either the D0XX or D1XX series of RAID set device numbers. Which series the hdisk are mapped to depends on the comparison between the S value of the CCL and the hdisk in question. The host adapter string of the location code of the CCL hdisk must match the target hdisk adapter string before the series question can be answered. For the target hdisk, any of these whose S value is the same as the CCL S value is associated with the D0XX. The rest of the hdisk for the same host adapter are D1XX series. The L value in the AIX device location code resolves the XX value for the appropriate D0XX or D1XX to which the hdisk are mapped.

The following is an example for identifying the RAID set units in a single host configuration using testvg as a hypothetical volume group.

```
# lsvg -p testvg
testvg:
PV_NAMEPV STATETOTAL PPsFREE PPsFREE DISTRIBUTION
hdisk5active54222  00..00..00..00..22
hdisk6active54223  00..00..00..00..23

# lsdev -Cc disk
hdisk0  Available 10-60-00-8,0 16 Bit SCSI Disk Drive
hdisk1  Available 20-60-00-4,0 Other SCSI Disk Drive
hdisk2  Available 30-60-00-0,0 DEC HSG80 RAID Array
hdisk3  Available 30-60-00-1,0 DEC HSG80 Command Console LUN
hdisk4  Available 30-60-00-0,1 DEC HSG80 RAID Array
hdisk5 Available 30-60-00-0,2 DEC HSG80 RAID Array
hdisk6 Available 30-60-00-0,1 DEC HSG80 RAID Array
hdisk7  Available 30-60-00-1,2 DEC HSG80 RAID Array
hdisk8  Available 30-60-00-1,3 DEC HSG80 RAID Array
hdisk9  Available 30-60-00-1,4 DEC HSG80 RAID Array

#
```

Note: The target hdisks from `lsvg` command are hdisk5 and hdisk6.

Note: The CCL (hdisk 3) is the same host adapter as hdisk 5 and hdisk6.

Note: The S value for hdisk6 is the same as the CCL. The L value for hdisk6 is 1 therefore the RAID set unit number is D001.

Note: Since the S value is not the same as the CCL and L value is 2, the RIAD set unit number for hdisk5 is D102.

For dual host connection configuration, the S value for the hdisk location code is zero. However, any hdisks, associated with the D0XX and D1XX series storage system units, is distinguished by their host adapter associations. As a consequence, the CCL, by its host designation, determines the series.

Refer to the following dual host example:

```
# lsdev -Cc disk
hdisk0  Available 10-60-00-8,0 16 Bit SCSI Disk Drive
hdisk1  Available 20-60-00-4,0 Other SCSI Disk Drive
hdisk2  Available 30-60-00-0,0 DEC HSG80 RAID Array
hdisk3  Available 40-60-00-0,0 DEC HSG80 Command Console LUN
hdisk4  Available 30-60-00-0,1 DEC HSG80 RAID Array
hdisk5 Available 30-60-00-0,2 DEC HSG80 RAID Array
hdisk6 Available 40-60-00-0,1 DEC HSG80 RAID Array
hdisk7  Available 40-60-00-0,2 DEC HSG80 RAID Array
hdisk8  Available 40-60-00-0,3 DEC HSG80 RAID Array
hdisk9  Available 40-60-00-0,4 DEC HSG80 RAID Array

#
```

Note: The target hdisks from `1svg` command are hdisk5 and hdisk6.

Note: The CCL (hdisk 3) is the same host adapter as hdisk6. The L value for hdisk6 is 1 therefor the RAID set unit number is D001.

Note: Since hdisk5 host adapter string does not match the CCL and the L value is 2, the RAID set unit number for hdisk5 is D102.

From this result, D001 and D102 are the target drives for copy in the storage system.

4. If the copy is to be a Snapshot, skip to Step 5, otherwise follow the Clone procedure described in this guide. Repeat the procedure for each target drive. As mentioned before, the hdisk associated with the storage system drives and targeted for Clone copy can remain available and logical volumes mounted in AIX. However, no AIX operation that results in a write to the drives should be allowed. When the Clone is complete, go to Step 6.
5. For Snapshot, the HSG80 controller must loaded with “S” type firmware before this function can succeed. Select and initialize the disk devices in the storage system having the same capacity as the target drives. Create a one-disk mirror container for each selected drive. Using a terminal attached to the HSG80 controller serial interface, add the drives as Snapshot units using the following CLI command construct:

```
ADD SNAPSHOT <New Unit Number> <Mirror Container Name> <Target Disk Unit>
```

6. Depending on the storage system configuration with the host or hosts, it may be necessary to control or activate host access to the new drive. Refer to the section, "[Restricting Host Access in Multiple-Bus Failover Mode](#)", page 55. Once this has been addressed, the new drives should be available for AIX use.
7. Run `cfgmgr` on the AIX system after which a new set of hdisks equal to the number of target hdisks should become available.
8. The new hdisks can be verified as copies of the originals by examining the respective hdisk physical volume IDs (PVID). To do this, issue the `lspv` AIX command.

The following is an example:

```
# lspv
hdisk0 000518660c11216F rootvg
hdisk1 003b5075af8b1c4e datavg
hdisk2 003b50756c0fc93e datavg
hdisk3 003b5075af8b1c4e none
hdisk4 003b50756c0fc93e none
```

hdisk3 and hdisk4 are the copied drives created by the storage system disk copy procedure (Clone or Snapshot) that also became available after the AIX cfgmgr was invoked. These new physical volumes are not associated with a volume group and the PVID of both drives match those respectively in hdisk1 and hdisk2.

9. Change the new hdisks PVIDs.

In the example in Step 8, hdisk1 and hdisk2 are part of an active volume group, however, AIX won't let hdisk3 and hdisk4 to be imported because the PVIDs are already known by the system. Changing the PVIDs for each new hdisk can resolve this issue using the following command for each new hdisk copy:

```
chdev -l <New hdisk> -a pv=clear; chdev -l <New hdisk> -a pv=yes
```

10. Use the recreatevg AIX command to import and activate the new hdisks into a volume group.

Although Step 9 solves the PVID problem in importing the new hdisks as a volume group, there are still conflicts with the logical volume found on the new hdisks and those from the original. `recreatevg` is an undocumented AIX command, available only in V4.3.3, that solves this logical volume issue. It also imports and activates the volume group. Use the following command:

```
recreatevg -y <New Volume Group Name> <One of the new hdisks set>
```

`recreatevg` also creates a mount point, `/fs`, under which all mount points from the original volume group are created, i.e., if the original mount is `/disk1`, for example, then the path for the new volume group path is `/fs/disk1`.

11. Mount the new volume group's logical volumes, after which the new logical volumes should be ready for use.

Moving StorageSets

You can move a storageSet from one subsystem to another without destroying its data. You also can follow the steps in this section to move a storageSet to a new location within the same subsystem.



Caution: Move only normal storageSets. Do not move storageSets that are reconstructing or reduced, or data corruption results.

See the release notes for the version of your controller software for information on which drives can be supported.



Caution: Never initialize any container or this procedure does not protect data in the storageSet.

Use the following procedure to move a storageSet, while maintaining the data the storageSet contains:

1. Show the details for the storageSet you want to move. Use the following command:

```
SHOW STORAGESET-NAME
```

2. Label each member with its name and PTL location.

If you do not have a storageSet map for your subsystem, you can enter the LOCATE command for each member to find its PTL location. Use the following command:

```
LOCATE DISK-NAME
```

To cancel the locate command, enter the following:

```
LOCATE CANCEL
```

3. Delete the unit number shown in the “Used by” column of the SHOW *storageSet-name* command. Use the following syntax:

```
DELETE UNIT-NUMBER
```

4. Delete the storageSet shown in the “Name” column of the SHOW *storageSet-name* command. Use the following syntax:

```
DELETE STORAGESET-NAME
```

5. Delete each disk drive, one at a time, that the storageset contained. Use the following syntax:

```
DELETE DISK-NAME
```

```
DELETE DISK-NAME
```

```
DELETE DISK-NAME
```

6. Remove the disk drives and move them to their new PTL locations.
7. Again add each disk drive to the controller's list of valid devices. Use the following syntax:

```
ADD DISK DISK-NAME PTL-LOCATION
```

```
ADD DISK DISK-NAME PTL-LOCATION
```

```
ADD DISK DISK-NAME PTL-LOCATION
```

8. Recreate the storageset by adding its name to the controller's list of valid storagesets and by specifying the disk drives it contains. (Although you have to recreate the storageset from its original disks, you do not have to add the storagesets in their original order.) Use the following syntax to recreate the storageset:

```
ADD STORAGESET-NAME DISK-NAME DISK-NAME
```

9. Represent the storageset to the host by giving it a unit number the host can recognize. You can use the original unit number or create a new one. Use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESSET-NAME
```

The following example moves unit D100 to another cabinet. D100 is the RAIDset RAID99 that consists of members DISK10000, DISK20000, and DISK10100.

Old cabinet

```
DELETE D100  
DELETE RAID99  
DELETE DISK10000  
DELETE DISK10100  
DELETE DISK20000  
DELETE DISK20100
```

New cabinet

```
ADD DISK DISK10000
ADD DISK DISK10100
ADD DISK DISK20000
ADD DISK DISK20100
ADD RAIDSET RAID99 DISK10000 DISK10100 DISK20000 DISK20100
ADD UNIT D100 RAID99
```


Subsystem Profile Templates

A

This appendix contains storageset profiles to copy and use to create your profiles. It also contains an enclosure template to use to help keep track of the location of devices and storagesets in your shelves. Four (4) templates are needed for the subsystem.

Note: The storage map templates for the Model 4310R and Model 4214R or 4314R reflect the physical location of the disk enclosures in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.

- "Storageset Profile", page 196
- "Storage Map Template 1 for the BA370 Enclosure", page 198
- "Storage Map Template 2 for the Second BA370 Enclosure", page 199
- "Storage Map Template 3 for the Third BA370 Enclosure", page 200
- "Storage Map Template 4 for the Model 4214R Disk Enclosure", page 201
- "Storage Map Template 5 for the Model 4254 Disk Enclosure", page 203
- "Storage Map Template 6 for the Model 4310R Disk Enclosure", page 205
- "Storage Map Template 7 for the Model 4350R Disk Enclosure", page 208
- "Storage Map Template 8 for the Model 4314R Disk Enclosure", page 210
- "Storage Map Template 9 for the Model 4354R Disk Enclosure", page 213

StorageSet Profile

Type of StorageSet:

<input type="checkbox"/> Mirrorset	<input checked="" type="checkbox"/> RAIDset	<input type="checkbox"/> Stripeset	<input type="checkbox"/> Striped	<input type="checkbox"/> JBOD
Mirrorset				

StorageSet Name**Disk Drives****Unit Number**

Partitions:

| Unit # |
--------	--------	--------	--------	--------	--------	--------	--------

RAIDset Switches:

Reconstruction Policy	Reduced Membership	Replacement Policy
<input type="checkbox"/> Normal (default)	<input type="checkbox"/> No (default)	<input type="checkbox"/> Best performance (default)
<input type="checkbox"/> Fast	<input type="checkbox"/> Yes, missing:	<input type="checkbox"/> Best fit
		<input type="checkbox"/> None

Mirrorset Switches:

Replacement Policy	Copy Policy	Read Source
<input type="checkbox"/> Best performance (default)	<input type="checkbox"/> Normal (default)	<input type="checkbox"/> Least busy (default)
<input type="checkbox"/> Best fit	<input type="checkbox"/> Fast	<input type="checkbox"/> Round robin
<input type="checkbox"/> None		<input type="checkbox"/> Disk drive:

Initialize Switches:

Chunk size	Save Configuration	Metadata
<input type="checkbox"/> Automatic (default)	<input type="checkbox"/> No (default)	<input type="checkbox"/> Destroy (default)
<input type="checkbox"/> 64 blocks	<input type="checkbox"/> Yes	<input type="checkbox"/> Retain
<input type="checkbox"/> 128 blocks		
<input type="checkbox"/> 256 blocks		
<input type="checkbox"/> Other:		

Unit Switches:

Caching	Access by following hosts enabled
Read caching _____	_____
Read-ahead caching _____	_____
Write-back caching _____	_____
Write-through caching _____	_____

Storage Map Template 1 for the BA370 Enclosure

Use this template for:

- BA370 single-enclosure subsystems
- first enclosure of multiple BA370 enclosure subsystems

	Port						
	1	2	3	4	5	6	
Power Supply	D10300	D20300	D30300	D40300	D50300	D60300	Power Supply
Power Supply	D10200	D20200	D30200	D40200	D50200	D60200	Power Supply
Power Supply	D10100	D20100	D30100	D40100	D50100	D60100	Power Supply
Power Supply	D10000	D20000	D30000	D40000	D50000	D60000	Power Supply

Targets
3
2
1
0

Storage Map Template 2 for the Second BA370 Enclosure

Use this template for the second enclosure of multiple BA370 enclosure subsystems.

	Port						
	1	2	3	4	5	6	
Power Supply							Power Supply
	D11100	D21100	D31100	D41100	D51100	D61100	
Power Supply							Power Supply
	D11000	D21000	D31000	D41000	D51000	D61000	
Power Supply							Power Supply
	D10900	D20900	D30900	D40900	D50900	D60900	
Power Supply							Power Supply
	D10800	D20800	D30800	D40800	D50800	D60800	

Targets

11
10
9
8

Storage Map Template 3 for the Third BA370 Enclosure

Use this template for the third enclosure of multiple BA370 enclosure subsystems.

	Port						
	1	2	3	4	5	6	
Power Supply							Power Supply
	D11500	D21500	D31500	D41500	D51500	D61500	
Power Supply							Power Supply
	D11400	D21400	D31400	D41400	D51400	D61400	
Power Supply							Power Supply
	D11300	D21300	D31300	D41300	D51300	D61300	
Power Supply							Power Supply
	D11200	D21200	D31200	D41200	D51200	D61200	

Targets

15
14
13
12

Storage Map Template 4 for the Model 4214R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4214R disk enclosure (single-bus). You can have up to six Model 4214R disk enclosures per controller shelf.

Model 4214R Disk Enclosure Shelf 1 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500

Model 4214R Disk Enclosure Shelf 2 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500

Model 4214R Disk Enclosure Shelf 3 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

Storage Map Template 5 for the Model 4254 Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4254 disk enclosure (dual-bus). You can have up to three Model 4254 disk enclosures per controller shelf.

Model 4254 Disk Enclosure Shelf 1 (Dual-bus)														
Bay	Bus A							Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800

Model 4254 Disk Enclosure Shelf 2 (Dual-bus)														
Bay	Bus A							Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

continued on the following page

continued from previous page

Model 4254 Disk Enclosure Shelf 3 (Dual-bus)														
Bay	Bus A							Bus B						
	1	2	3	4	5	6	7	8	9	1	0	1	1	1
SCSI ID	0 0	0 1	0 2	0 3	0 4	0 5	0 8	0 0	0 1	0 2	0 3	0 4	0 5	0 8
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

Storage Map Template 6 for the Model 4310R Disk Enclosure

Use this template for a subsystem with a six-shelf Model 4310R disk enclosure (single-bus). You can have up to six Model 4310R disk enclosures per controller shelf.

Model 4310R Disk Enclosure Shelf 6 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200

Model 4310R Disk Enclosure Shelf 5 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200

Model 4310R Disk Enclosure Shelf 4 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

Model 4310R Disk Enclosure Shelf 1 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk11000	Disk11100	Disk11200

Model 4310R Disk Enclosure Shelf 2 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk21000	Disk21100	Disk21200

Model 4310R Disk Enclosure Shelf 3 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk31000	Disk31100	Disk31200

Storage Map Template 7 for the Model 4350R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4350R disk enclosure (single-bus). You can have up to three Model 4350R disk enclosures per controller shelf.

Model 4350R Disk Enclosure Shelf 6 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
Model 4350R Disk Enclosure Shelf 5 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200

Model 4350R Disk Enclosure Shelf 4 (Single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

Storage Map Template 8 for the Model 4314R Disk Enclosure

Use this template for a subsystem with a six-shelf Model 4314R disk enclosure. You can have a maximum of six Model 4314R disk enclosures with each Model 2200 controller enclosure.

Model 4314R Disk Enclosure Shelf 6 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk60900	Disk61000	Disk61100	Disk61200	Disk61300	Disk61400	Disk61500

Model 4314R Disk Enclosure Shelf 5 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk50900	Disk51000	Disk51100	Disk51200	Disk51300	Disk51500	Disk51500

Model 4314R Disk Enclosure Shelf 4 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk41100	Disk41200	Disk41300	Disk41400	Disk41500
Model 4314R Disk Enclosure Shelf 1 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500
Model 4314R Disk Enclosure Shelf 2 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500

Model 4314R Disk Enclosure Shelf 3 (Single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

Storage Map Template 9 for the Model 4354R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4354R disk enclosure (dual-bus). You can have up to three Model 4354R disk enclosures per controller shelf.

Model 4354R Disk Enclosure Shelf 1 (Dual-bus)														
Bay	SCSI Bus A							SCSI Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800
Model 4354R Disk Enclosure Shelf 2 (Dual-bus)														
Bay	SCSI Bus A							SCSI Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

Model 4354R Disk Enclosure Shelf 3 (Dual-bus)														
Bay	SCSI Bus A							SCSI Bus B						
SCSI ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Disk50000	00	01	02	03	04	05	08	00	01	02	03	04	05	08
Disk50100														
Disk50200														
Disk50300														
Disk50400														
Disk50500														
Disk50800														
Disk60000														
Disk60100														
Disk60200														
Disk60300														
Disk60400														
Disk60500														
Disk60800														

Installing, Configuring, and Removing the Client

B

The following information is included in this appendix:

- "Why Install the Client?", page 216
- "Before You Install the Client", page 217
- "Installing the Client", page 218
- "Installing the Integration Patch", page 219
- "Troubleshooting Client Installation", page 222
- "Adding Storage Subsystem and its Host to Navigation Tree", page 224
- "Removing Command Console Client", page 226
- "Where to Find Additional Information", page 228

Why Install the Client?

The Client monitors and manages a storage subsystem by performing the following tasks:

- Create mirrored device group (RAID 1)
- Create striped device group (RAID 0)
- Create striped mirrored device group (RAID 0+1)
- Create striped parity device group (3/5)
- Create an individual device (JBOD)
- Monitor many subsystems at once
- Set up pager notification

Before You Install the Client

1. Verify that you are logged into an account that is a member of the administrator group.
2. Check the software product description that came with the software for a list of supported hardware.
3. Verify that you have the SNMP service installed on the computer. SNMP must be installed on the computer for this software to work properly. The Client software uses SNMP to receive traps from the Agent. The SNMP service is available on the Windows NT, Windows 2000, and Windows Server 2003 installation web kit. To verify that you have the SNMP service:
 - For Windows NT, double-click Services in **Start > Settings > Control Panel**. The entry for SNMP is shown in this window. If you install the SNMP service and you already have Windows NT Service Pack 6A on the computer, reinstall the service pack after installing the SNMP service.
 - For Windows 2000 and Windows Server 2003, click **Start > Settings > Control Panel > Administrative Tools > Services**. The entry for SNMP is shown in the Services window.
4. Read the release notes.
5. Read "[Installing the Integration Patch](#)", page 219, in this appendix.
6. If you have the Command Console Client open, exit the Command Console Client.
7. If you have Command Console Client V1.1b or earlier, remove the program with the Windows **Add/Remove Programs** utility.
8. If you have a previous version of Command Console, you can save the Navigation Tree configuration by copying the `SWCC2.MDB` file to another directory. After you have installed the product, move `SWCC2.MDB` to the directory to which you installed SWCC.
9. Install the HS-Series Agent. For more information, see Chapter 4.
10. Download the same IBM AIX Platform Kit for HSG80 into an empty directory on your Windows system as was done for your AIX IBM server.
11. Untar the kit into this directory.

Note: Utilities like WinZip can be used to untar the kit.

Installing the Client

The following restriction should be observed when installing SWCC on Windows NT 4.0 Workstations.

If you select all of the applets during installation, the installation fails on the HSG60 applet and again on one of the HSG80 applets. The workaround is to install all of the applets you want except for the HSG60 applet and the HSG80 ACS 8.5 applet. You can then return to the setup program and install the one that you need.

1. In a SAN environment where you would need both HSG60 and HSG80 subsystems, HP recommends you install both, but one at a time. This problem is not seen under Windows NT 4.0 Server.
2. Set your Windows Explorer to the directory of your extracted AIX Solution Kit and run `autoinstall.exe`.

A dialog box is displayed.

3. One of the items in the dialog box should say **SWCC Client Software** and has a button that says **INSTALL** next to it. Click on the button to start the SWCC Client installation procedure.
4. Select the **HSG80 Controller for ACS85** or newer menu option to properly install SWCC Client, and click **Next**.

If this method does not work, run the `setup.exe` program.

Note: If the computer does not find a previous installation, it installs the SWCC Navigation Window and the CLI Window.

5. Follow the instructions on the screen. After you install the software, the Asynchronous Event Service (AES) starts. AES is a service that runs in the background. It collects and passes traps from the subsystems to the Navigation Tree and to individual pagers (for example, to show that a disk has failed). AES needs to be running for the Client system to receive updates.

Note: For more information on AES, see *HP StorageWorks Command Console V2.5 User Guide*.

Installing the Integration Patch

The integration patch determines which version of firmware the controller is using and launches the appropriate StorageWorks Command Console (SWCC) Storage Window within Insight Manager (CIM) V4.23.

Should I Install the Integration Patch?

Install this patch if your HSG80 controller uses ACS 8.6 or later. This patch enables you to use the controller's SWCC Storage Window within CIM to monitor and manage the controller.

Installing the Integration Patch

Perform the following steps to install the integration patch.

1. Verify that you have installed the HSG80 Storage Window for ACS 8.6 or later in the **Add/Remove Programs** in the Windows Control Panel. The HSG80 Storage Window for ACS 8.6 or later is needed to open the correct Storage Window for your version of the firmware.
2. Verify that you have installed HSG80 Storage Window V2.1 in the **Add/Remove Programs** (HP StorageWorks HSG80 V2.1) in the Windows Control Panel. The HSG80 Storage Window V2.1 is required to run the integration patch.
3. Verify that you have installed CIM V4.23.

Install the integration patch by double-clicking on `SWCC\Client\HSG80shim\setup.exe` in the AIX kit directory. The patch is installed in the same location as the original SWCC installation.

Caution: Do not remove the HSG80 Client from your computer. If you remove the HSG80 Client, you no longer are able to access its Storage Window.

Caution: If you remove the integration patch, HSG80 Storage Window V2.1 no longer works and you need to reinstall HSG80 Storage Window V2.1. The integration patch uses some of the same files as the HSG80 Storage Window V2.1.

Integrating Controller's SWCC Storage Window with CIM

You can open the controller's Storage Window from within the Windows-based CIM V4.23 by doing the following:

1. Verify that you have installed the following by looking in **Add/Remove Programs** in Control Panel:
 - The HSG80 Storage Window for ACS V8.6 or later (Required to open the correct Storage Window for your firmware).
 - The HSG80 Storage Window V2.1 (HP StorageWorks HSG80 V2.1) The CIM integration patch uses files in this program.
 - CIM V4.23.
 - CIM integration patch (HSG80 Insight Manager Shim).
2. Verify that you have installed the CIM Agent and the StorageWorks Command Console HS-Series Agent on the same computer.
3. Add the name of the Client system that has CIM to the Agent's list of Client system entries and choose SNMP as a notification scheme.
4. Open Insight Manager.
5. To open the Server window, click on the device you want to observe in the CIM Navigation window.
6. Click on the **Mass Storage** button in the Server window. The **CIM Navigation Tree** opens.
7. Click on the + symbol next to **RAID Storage System**. The Navigation Tree expands to display a listing called **Storage System Information**.
8. Double-click **Storage System Information**. You are given the status of the system.
9. Click **Launch**. The controller's Storage Window opens.

"Insight Manager Unable to Find Controller's Storage Window"

If you installed Insight Manager before SWCC, Insight Manager is unable to find the controller's Storage Window. To find the controller's Storage Window, perform the following procedure:

1. Double-click the **Insight Agents** icon (**Start > Settings > Control Panel**). A window is displayed showing you the active and inactive Agents under the **Services** tab.

2. Highlight the entry for **Fibre Array Information** and click **Add**. The Fibre Array Information entry is moved from **Inactive Agents** to **Active Agents**.

Troubleshooting Client Installation

This section provides information on how to resolve some of the problems that may occur when installing the Client software:

- Invalid Network Port Assignments During Installation
- “There is no disk in the drive” Message

Invalid Network Port Assignments During Installation

SWCC Clients and Agents communicate by using sockets. The SWCC installation attempts to add entries into each system list of services (services file or for UCX, the local services database). If the SWCC installation finds an entry in the local services file with the same name as the one it wants to add, it assumes the one in the file is correct.

The SWCC installation may display a message, stating that it cannot upgrade the services file. This happens if it finds an entry in the local services file with the same number as the one it wants to add, but with a different name. In that case, appropriate port numbers must be obtained for the network and added manually to the services file.

There are two default port numbers, one for Command Console (4998) and the other for the device-specific Agent and Client software, such as the Fibre Channel Interconnect Client and Agent (4989). There are two exceptions. The following software has two default port numbers:

- The KZPCC Agent and Client, (4991 and 4985)
- The RA200 Agent and Client, (4997 and 4995)

If the Network Information Services (NIS) are being used to provide named port lookup services, contact the network administrator to add the correct ports.

The following shows how the network port assignments display in the services file:

spgui	4998/tcp	#Command Console
ccdevmgt	4993/tcp	#Device Management Client and Agent
kzpccconnectport	4991/tcp	#KZPCC Client and Agent
kzpccdiscoveryport	4985/tcp	#KZPCC Client and Agent

ccfabric	4989/tcp	#Fibre Channel Interconnect Agent
spagent	4999/tcp	#HS-Series Client and Agent
spagent3	4994/tcp	#HSZ22 Client and Agent
ccagent	4997/tcp	#RA200 Client and Agent
spagent2	4995/tcp	#RA200 Client and Agent

"There is no disk in the drive" Message

When you install the Command Console Client, the software checks the shortcuts on the desktop and in the Start menu. The installation checks the shortcuts of all users for that computer, even if they are not currently logged on. You may receive an error message if any of these shortcuts point to empty floppy drives, empty CD-ROM drives, or missing removable disks. Do one of the following:

- Ignore the error message by clicking **Ignore**.
- Replace the removable disks, and place a disk in the floppy drive and a CD-ROM in the CD-ROM drive. Then, click **Retry**.

Adding Storage Subsystem and its Host to Navigation Tree

The Navigation Tree enables you to manage storage over the network by using the Storage Window. If you plan to use pager notification, you must add the storage subsystem to the Navigation Tree.

1. Verify that you have properly installed and configured the HS-Series Agent on the storage subsystem host.
2. Click **Start > Programs > Command Console > StorageWorks Command Console**.

The Client displays the Navigation Window. The Navigation Window lets you monitor and manage many storage subsystems over the network.




Figure 45: Navigation window

3. Click **File > Add System**.

The Add System window is displayed.

4. Type the host name or its TCP/IP address and click **Apply**.
5. Click **Close**.




Figure 46: Navigation window showing storage host system "Atlanta"

6. Click the plus sign to expand the host icon. When expanded, the Navigation Window displays an icon for the storage subsystem. To access the Storage Window for the subsystem, double-click the Storage Window icon.



Figure 47: Navigation window showing expanded "Atlanta" host icon

Note: You can create virtual disks by using the Storage Window. For more information on the Storage Window, refer to *HP StorageWorks Command Console V2.5 User Guide*.

Removing Command Console Client

Before you remove the Command Console Client from the computer, remove AES. This prevents the system from reporting that a service failed to start every time the system is restarted. Steps 2 through 5 describe how to remove the Command Console Client.

Note: When you remove the Command Console Client, the `SWCC2.MDB` file is deleted. This file contains the Navigation Tree configuration. If you want to save this information, move the file to another directory.

1. Click **Start > Programs > StorageWorks Command Console** and change to the directory where you installed the Command Console Client.
2. Enter the following command:
`C:\Program Files\Compaq\SWCC> AsyncEventService -remove`
3. Do one of the following:
 - On Windows NT 4.0, click **Start > Settings > Control Panel**, and double-click the **Add/Remove Programs** icon in the Control Panel. The **Add/Remove Program Properties** window is displayed.
 - On Windows 2000 and Windows Server 2003, click **Start > Settings > Control Panel > Add/Remove Programs**. The **Add/Remove Programs** window is displayed.
4. Select **Command Console** in the window.
5. Do one of the following:
 - On Windows NT 4.0, click **Add/Remove**.
 - On Windows 2000 and Windows Server 2003, click **Change/Remove**.
6. Follow the instructions on the screen.

Note: This procedure removes only the Command Console Client (SWCC Navigation Window). You can remove the HSG80 Client by using the Add/Remove program.

Where to Find Additional Information

You can find additional information about SWCC by referring to the online Help and to *HP StorageWorks Command Console V2.5 User Guide*.

About the User Guide

HP StorageWorks Command Console V2.5 User Guide contains additional information on how to use SWCC. Some of the topics in the user guide are the following:

- About AES
- Adding Devices
- Adding Virtual Disks
- Setting Up Pager Notification
- How to Integrate SWCC with Insight Manager
- Troubleshooting Information

About the Online Help

Most of the information about the Client is provided in the online Help. Online Help is provided in two places:

- Navigation Window—Online Help provides information on pager notification and a tour of the Command Console Client, in addition to information on how to add a system to the Navigation Tree.
- Storage Window—Online Help provides detailed information about the Storage Window, such as how to create virtual disks.

HACMP Notes for AIX



Supported Versions of HACMP

HACMP is the acronym for High Availability Cluster Multi-Processing.

Only HACMP/ES 4.4.1.x and 4.5 are supported with the solution software.

Installing Solution Software for AIX V8.8 on HACMP Cluster

The installation procedures are the same as those for a standalone node. The solution software must be loaded on each server in the cluster as if it were a standalone node.

Uninstalling Solution Software for AIX on HACMP Cluster

The uninstall procedure is the same as for a standalone node. It is important to note that you should remove any volume groups that are associated with any HSG80 RAID Array before you execute either `cpq_hsg_uninstall` or the `UNINSTALL` utility on the solution software CD-ROM. The solution software must be uninstalled on each node in the cluster.

Adding a Virtual Disk After IPL to HACMP Cluster

The procedure for adding a volume group in an HACMP cluster is different than adding one on a stand-alone server. The difference here is that you must create a shared volume group between the two servers that can be managed by HACMP. The following is a procedure to create a shared volume group on two servers in a cluster.

Creating a Volume Group on a Virtual Disk

This procedure was done on a two node cluster, brutus and miranda, running HACMP/ES 4.4.1. This procedure does not supersede the procedure outlined in the HACMP documentation on creating a shared volume group. These steps are

just meant as guide to you. Consult your HACMP documentation for any changes on creating a shared volume group. Note that using C-SPOC for HACMP does not work for HSG80 RAID Array devices.

1. Show a list of devices on each node in the cluster to see what device you want to create a shared volume group on by issuing the command `lshsg` on each node.

Output from brutus:

```
hdisk3 Available 37-08-00-4,0 DEC HSG80 Command Console LUN  
hdisk5 Available 37-08-00-4,1 DEC HSG80 RAID Array  
hdisk6 Available 3A-08-00-4,0 DEC HSG80 Command Console LUN
```

Output from miranda:

```
hdisk3 Available 17-08-00-4,0 DEC HSG80 Command Console LUN  
hdisk4 Available 17-08-00-4,1 DEC HSG80 RAID Array  
hdisk5 Available 1A-08-00-4,0 DEC HSG80 Command Console LUN
```

Here, we use `hdisk5` from node brutus and `hdisk4` from node miranda. These are actually the same disk being presented to different nodes.

2. Find out the Physical Volume Identifier (PVID) using the command `lspv` on each node.

Output from brutus:

hdisk0	0002997fb3fcfab5	none
hdisk1	0002997f32d9425d	rootvg
hdisk2	none	none
hdisk3	none	none
hdisk4	none	none
hdisk5	0002997f9d52748a	none
hdisk6	none	none

Output from miranda:

hdisk0	0000981f0ebaf585	none
hdisk1	0000981f6f1f38bb	none
hdisk2	0000981f6c12e583	rootvg
hdisk3	none	none
hdisk4	0002997f9d52748a	none
hdisk5	none	none

hdisk6	none	none
--------	------	------

There is no PVID for our disk on either node, so we assign a PVID on one node and then get the other node to use the same PVID.

From here, you want to pick one node to do the entire configuration on and then use the other node to simply import the needed information. In this case we use node brutus as our main node and miranda as our other node.

3. On node brutus, execute the command

```
chdev -a pv=yes -l hdisk5
```

4. On node miranda, execute the command

```
chdev -a pv=yes -l hdisk4
```

5. Using the `lspv` command, examine the PVID on each disk to make sure that they are the same.

Output from brutus:

hdisk0	0002997fb3fcfab5	none
hdisk1	0002997f32d9425d	rootvg
hdisk2	none	none
hdisk3	none	none
hdisk4	none	none
hdisk5	none	none
hdisk6	none	none

Output from miranda:

hdisk0	0000981f0ebaf585	none
hdisk1	0000981f6f1f38bb	none
hdisk2	0000981f6c12e583	rootvg
hdisk3	none	none
hdisk4	none	none
hdisk5	none	none
hdisk6	none	none

6. Execute the command `lvlstmajor` on each node in the cluster to get the next available major number to use for a volume group:

Output from brutus:

```
45...
```

Output from miranda:

44...

You are going to want to use the greater of the two numbers, in this case 45.

7. Create the volume group on node brutus (all commands are on brutus until otherwise specified) using the command:

```
mkvg -f -V 45 -s 32 -y sharedvg hdisk5
```

Note: The arguments to `mkvg` can and do vary. Use arguments that are appropriate to the size and purpose of your device.

8. Execute the command `lsvg sharedvg` to get the number of Free Physical Partitions.

This number can be found in the “Free PPs:” field.

The output should look similar to the following:

```
VOLUME GROUP:    sharedvg
VG IDENTIFIER:  0002997f00004c00000000ec959cb382
VG STATE:        active PP SIZE: 32 megabytes
VG PERMISSION:   read/write TOTAL PPs:      31 (992 megabytes)
MAX LVs:         256 FREE PPs: 31 (992 megabytes)
LVs:             0 USED PPs:      0 (0 megabytes)
OPEN LVs:        0 QUORUM:          2
TOTAL PVs:       1 VG DESCRIPTORS: 2
STALE PVs:       0 STALE PPs:      0
ACTIVE PVs:      1 AUTO ON:        yes
MAX PPs per PV: 1016 MAX PVs:        32
LTG size:        128 kilobytes AUTO SYNC:     no
```

In this case, the number of free PPs is 31.

9. Create a logical volume using a command similar to

```
mklv -y sharedlv sharedvg 30
```

Note: We use the number of Free PPs - 1 because there needs to be 1 PP free to create a filesystem.

Note: The arguments to `mklv` can and do vary. Use arguments that are appropriate to the purpose of your device.

10. Create a filesystem on the logical volume now. In this case, we are going to create a filesystem that spans the whole disk. This is not mandatory. Use a command similar to:

```
crfs -v jfs -d sharedlv -m /sharedfs
```

Note: The arguments to `crfs` can and do vary. Use arguments that are appropriate to the purpose of the filesystem.

11. Quickly mount and unmount the filesystem at it's mount point to ensure that the filesystem has been correctly made.

```
mount /sharedfs  
umount /sharedfs
```

12. Change the volume group so that it does not autovaryon at boot time.

```
chvg -a n sharedvg
```

13. Deactivate the volume group so that it can be imported by the other node.

```
varyoffvg sharedvg
```

14. On node miranda (the rest of the commands apply to node miranda now), execute the `importvg` command to get the volume group definition from the device itself.

```
importvg -V 45 -y sharedvg hdisk4
```

15. Change the volume group so that it does no autovaryon at boot time.

```
chvg -a n sharedvg
```

16. Deactivate the volume group

```
varyoffvg sharedvg
```

The Virtual Disk is now ready to be added to your HACMP configuration.

glossary

This glossary defines terms pertaining to the ACS solution software. It is not a comprehensive glossary of computer terms.

8B/10B	A type of byte definition encoding and decoding to reduce errors in data transmission patented by the IBM Corporation. This process of encoding and decoding data for transmission has been adopted by ANSI.
adapter	A device that converts the protocol and hardware interface of one bus type into another without changing the function of the bus.
ACS	<i>See array controller software.</i>
AL_PA	<i>See arbitrated loop physical address.</i>
alias address	An AL_PA value recognized by an arbitrated loop port in addition to the assigned AL_PA.
ANSI	Pronounced “ann-see.” Acronym for the American National Standards Institute. An organization who develops standards used voluntarily by many manufacturers within the USA. ANSI is not a government agency.
arbitrate	A process of selecting one L_Port from a collection of several ports that request use of the arbitrated loop concurrently.
arbitrated loop	A loop type of topology where two or more ports can be interconnected, but only two ports at a time can communicate.
arbitrated loop physical address	Abbreviated AL_PA. A one-byte value used to identify a port in an Arbitrated Loop topology.
array controller	<i>See controller.</i>
array controller software	Abbreviated ACS. Software contained on a removable ROM program card that provides the operating system for the array controller.

association set	A group of remote copy sets that share selectable attributes for logging and failover. Members of an association set transition to the same state simultaneously. For example, if one association set member assumes the failsafe locked condition, then other members of the association set also assume the failsafe locked condition.
	An association set can also be used to share a log between a group of remote copy set members that require efficient use of the log space.
asynchronous	Pertaining to events that are scheduled as the result of a signal asking for the event; pertaining to that which is without any specified time relation. <i>See also</i> synchronous.
autospare	A controller feature that automatically replaces a failed disk drive. To aid the controller in automatically replacing failed disk drives, you can enable the AUTOSPARE switch for the failedset causing physically replaced disk drives to be automatically placed into the spareset. Also called “AUTONEWSPARE.”
bad block	A data block that contains a physical defect.
bad block replacement	Abbreviated BBR. A replacement routine that substitutes defect-free disk blocks for those found to have defects. This process takes place in the controller, transparent to the host.
backplane	The electronic printed circuit board into which you plug subsystem devices—for example, the SBB or power supply.
battery hysteresis	The ability of the software to allow write-block caching during the time a battery is charging, but only when a previous down time has not drained more than 50 percent of rated battery capacity.
BBR	<i>See</i> bad block replacement.
BIST	<i>See</i> built-in self-test.
bit	A single binary digit having a value of either 0 or 1. A bit is the smallest unit of data a computer can process.
block	Also called a sector. The smallest collection of consecutive bytes addressable on a disk drive. In integrated storage elements, a block contains 512 bytes of data, error codes, flags, and the block address header.
bootstrapping	A method used to bring a system or device into a defined state by means of its own action. For example, a machine routine whose first few instructions are enough to bring the rest of the routine into the computer from an input device.

built-in self-test	A diagnostic test performed by the array controller software on the controller policy processor.
byte	A binary character string made up of 8 bits operated on as a unit.
cache memory	A portion of memory used to accelerate read and write operations.
cache module	A fast storage buffer
CCL	CCL-Command Console LUN, a “SCSI Logical Unit Number” virtual-device used for communicating with Command Console Graphical User Interface (GUI) software.
channel	An interface that allows high speed transfer of large amounts of data. Another term for a SCSI bus. <i>See also</i> SCSI.
chunk	A block of data written by the host.
chunk size	The number of data blocks, assigned by a system administrator, written to the primary RAIDset or stripeset member before the remaining data blocks are written to the next RAIDset or stripeset member.
CLCP	An abbreviation for code-load code-patch utility. This utility is used to upgrade the controller and EMU software. It can also be used to patch the controller software.
CLI	See Command Line Interface.
coax	A two-conductor wire in which one conductor completely wraps the other with the two separated by insulation.
cold swap	A method of device replacement that requires the entire subsystem to be turned off before the device can be replaced. <i>See also</i> hot swap and warm swap.
command line interface	CLI. A command line entry utility used to interface with the HS-series controllers. CLI enables the configuration and monitoring of a storage subsystem through textual commands.
concat commands	Concat commands implement storageset expansion features.
configuration file	A file that contains a representation of a storage subsystem configuration.
container	1) Any entity that is capable of storing data, whether it is a physical device or a group of physical devices. (2) A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Stripesets and mirrorsets are examples of storageset containers the controller uses to create units.

controller	A hardware device that, with proprietary software, facilitates communications between a host and one or more devices organized in an array. The HSG80 family controllers are examples of array controllers.
copying	A state in which data to be copied to the mirrorset is inconsistent with other members of the mirrorset. <i>See also</i> normalizing.
copying member	Any member that joins the mirrorset after the mirrorset is created is regarded as a copying member. Once all the data from the normal member (or members) is copied to a normalizing or copying member, the copying member then becomes a normal member. <i>See also</i> normalizing member.
CSR	An acronym for control and status register.
DAEMON	Pronounced “demon.” A program usually associated with a UNIX systems that performs a utility (housekeeping or maintenance) function without being requested or even known of by the user. A daemon is a diagnostic and execution monitor.
data center cabinet	A generic reference to large subsystem cabinets, such as the cabinets in which HP StorageWorks components can be mounted.
data striping	The process of segmenting logically sequential data, such as a single file, so that segments can be written to multiple physical devices (usually disk drives) in a round-robin fashion. This technique is useful if the processor is capable of reading or writing data faster than a single disk can supply or accept the data. While data is being transferred from the first disk, the second disk can locate the next segment.
DDL	Dual data link. The ability to operate on the CI bus using both paths simultaneously to the same remote node.
device	<i>See</i> node and peripheral device.
differential I/O module	A 16-bit I/O module with SCSI bus converter circuitry for extending a differential SCSI bus. <i>See also</i> I/O module.
differential SCSI bus	A bus in which a signal level is determined by the potential difference between two wires. A differential bus is more robust and less subject to electrical noise than is a single-ended bus.
DIMM	Dual inline Memory Module.
dirty data	The write-back cached data that has not been written to storage media, even though the host operation processing the data has completed.
DMA	Direct Memory Access.

DOC	DWZZA-On-a-Chip. ASCSI bus extender chip used to connect a SCSI bus in an expansion cabinet to the corresponding SCSI bus in another cabinet (See DWZZA).
driver	A hardware device or a program that controls or regulates another device. For example, a device driver is a driver developed for a specific device that allows a computer to operate with the device, such as a printer or a disk drive.
dual-redundant configuration	A controller configuration consisting of two active controllers operating as a single controller. If one controller fails, the other controller assumes control of the failing controller devices.
dual-simplex	A communications protocol that allows simultaneous transmission in both directions in a link, usually with no flow control.
DUART	Dual universal asynchronous receiver and transmitter. An integrated circuit containing two serial, asynchronous transceiver circuits.
DWZZA	An HP StorageWorks SCSI bus signal converter used to connect 8-bit single-ended devices to hosts with 16-bit differential SCSI adapters. This converter extends the range of a single-ended SCSI cable to the limit of a differential SCSI cable.
DWZZB	An HP StorageWorks SCSI bus signal converter used to connect a variety of 16-bit single-ended devices to hosts with 16-bit differential SCSI adapters.
ECB	External cache battery. The unit that supplies backup power to the cache module in the event the primary power source fails or is interrupted.
ECC	Error checking and correction.
EDC	Error detection code.
EIA	The abbreviation for Electronic Industries Association. EIA is a standards organization specializing in the electrical and functional characteristics of interface equipment.
EMU	Environmental monitoring unit. A unit that provides increased protection against catastrophic failures. Some subsystem enclosures include an EMU which works with the controller to detect conditions such as failed power supplies, failed blowers, elevated temperatures, and external air sense faults. The EMU also controls certain cabinet hardware including DOC chips, alarms, and fan speeds.

ESD	Electrostatic discharge. The discharge of potentially harmful static electrical voltage as a result of improper grounding.
extended subsystem	A subsystem in which two cabinets are connected to the primary cabinet.
external cache battery	<i>See ECB.</i>
F_Port	A port in a fabric where an N_Port or NL_Port may attach.
fabric	A group of interconnections between ports that includes a fabric element.
failback	The process of restoring data access to the newly-restored controller in a dual-redundant controller configuration. See also failover.
failedset	A group of failed mirrorset or RAIDset devices automatically created by the controller.
failover	The process that takes place when one controller in a dual-redundant configuration assumes the workload of a failed companion controller. Failover continues until the failed controller is repaired or replaced.
	The ability for HSG80 controllers to transfer control from one controller to another in the event of a controller failure. This ensures uninterrupted operation. Use Transparent Failover mode for single HBA configurations. Use multiple-bus failover mode for Secure Path based configurations.
FCA	Fibre Channel Adapter
FC-AL	The Fibre Channel Arbitrated Loop standard. See <i>Fibre Channel</i> .
FC-ATM	ATM AAL5 over Fibre Channel
FC-FG	Fibre Channel Fabric Generic Requirements
FG-FP	Fibre Channel Framing Protocol (HIPPI on FC)
FC-GS-1	Fibre Channel Generic Services-1
FC-GS-2	Fibre Channel Generic Services-2
FC-IG	Fibre Channel Implementation Guide
FC-IE	Fibre Channel Link Encapsulation (ISO 8802.2)
FC-PH	The Fibre Channel Physical and Signaling standard.
FC-SB	Fibre Channel Single Byte Command Code Set
FC-SW	Fibre Channel Switched Topology and Switch Controls

FCC	Federal Communications Commission. The federal agency responsible for establishing standards and approving electronic devices within the United States.
FCC Class A	This certification label is on electronic devices that can only be used in a commercial environment within the United States.
FCC Class B	This certification label is on electronic devices that can be used in either a home or a commercial environment within the United States.
FCP	The mapping of SCSI-3 operations to Fibre Channel.
FDDI	Fiber Distributed Data Interface. An ANSI standard for 100 megabaud transmission over fiber optic cable.
FD SCSI	The fast, narrow, differential SCSI bus with an 8-bit data transfer rate of 10 MB/s. <i>See also</i> FWD SCSI and SCSI.
fiber	A fiber or optical strand. Spelled <i>fibre</i> in Fibre Channel.
fiber optic cable	A transmission medium designed to transmit digital signals in the form of pulses of light. Fiber optic cable is noted for its properties of electrical isolation and resistance to electrostatic contamination.
Fibre Channel	A high speed, high-bandwidth serial protocol for channels and networks that interconnect over twisted pair wires, coaxial cable or fiber optic cable. The Fibre Channel Switched (FC-SW) (fabric) offers up to 16 million ports with cable lengths of up to 10 kilometers. The Fibre Channel Arbitrated Loop (FC-AL) topology offers speeds of up to 100 Mbytes/seconds and up to 127 nodes, all connected in serial. In contrast to SCSI technology, Fibre Channel does not require ID switches or terminators. The FC-AL loop may be connected to a Fibre Channel fabric for connection to other nodes.
fibre channel topology	An interconnection scheme that allows multiple Fibre Channel ports to communicate with each other. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.
FL_Port	A port in a fabric where N_Port or an NL_Port may be connected.
flush	The act of writing dirty data from cache to a storage media.
FMU	Fault management utility.
forced errors	A data bit indicating a corresponding logical data block contains unrecoverable data.
frame	An invisible unit used to transfer information in Fibre Channel.

FRU	Field replaceable unit. A hardware component that can be replaced at the customer location by service personnel or qualified customer service personnel.
FRUTIL	Field Replacement utility.
full duplex (n)	A communications system in which there is a capability for 2-way transmission and acceptance between two sites at the same time.
full duplex (adj)	Pertaining to a communications method in which data can be transmitted and received at the same time.
FWD SCSI	A fast, wide, differential SCSI bus with a maximum 16-bit data transfer rate of 20 MB/s. <i>See also</i> SCSI and FD SCSI.
GBIC	Gigabit Interface Converter. GBICs convert electrical signals to optical signals (and vice-versa.) They are inserted into the ports of the Fibre Channel switch and hold the Fibre Channel cables.
GLM	Gigabit link module
giga	A prefix indicating a billion (10^9) units, as in gigabaud or gigabyte.
gigabaud	An encoded bit transmission rate of one billion (10^9) bits per second.
gigabyte	A value normally associated with a disk drives storage capacity, meaning a billion (10^9) bytes. The decimal value 1024 is usually used for one thousand.
half-duplex (adj)	Pertaining to a communications system in which data can be either transmitted or received but only in one direction at one time.
hard address	The AL_PA which an NL_Port attempts to acquire during loop initialization.
heterogeneous host support	Also called <i>noncooperating host support</i> .
HIPPI-FC	Fibre Channel over HIPPI
host	The primary or controlling computer to which a storage subsystem is attached.
host adapter	A device that connects a host system to a SCSI bus. The host adapter usually performs the lowest layers of the SCSI protocol. This function may be logically and physically integrated into the host system.
HBA	Host bus adapter

host compatibility mode	A setting used by the controller to provide optimal controller performance with specific operating systems. This improves the controller performance and compatibility with the specified operating system.
hot disks	A disk containing multiple hot spots. Hot disks occur when the workload is poorly distributed across storage devices which prevents optimum subsystem performance. <i>See also</i> hot spots.
hot spots	A portion of a disk drive frequently accessed by the host. Because the data being accessed is concentrated in one area, rather than spread across an array of disks providing parallel access, I/O performance is significantly reduced. <i>See also</i> hot disks.
hot swap or hot-pluggable	A method of device replacement that allows normal I/O activity on a device bus to remain active during device removal and insertion. The device being removed or inserted is the only device that cannot perform operations during this process. <i>See also</i> cold swap and warm swap.
hub	A device (concentrator) which performs some or all of the following functions:
	<ul style="list-style-type: none">■ Automatic insertion of operational loop devices without disrupting the existing configuration.■ Automatic removal of failed loop devices without impacting the existing configuration.■ Provides a centralized (star) wiring configuration and maintenance point.■ Provides central monitoring and management.
IBR	Initial Boot Record.
ILF	Illegal function.
INIT	Initialize input and output.
initiator	A SCSI device that requests an I/O process to be performed by another SCSI device, namely, the SCSI target. The controller is the initiator on the device bus. The host is the initiator on the host bus.
instance code	A four-byte value displayed in most text error messages and issued by the controller when a subsystem error occurs. The instance code indicates when during software processing the error was detected.

interface	A set of protocols used between components, such as cables, connectors, and signal levels.
I/O	Refers to input and output functions.
I/O driver	The set of code in the kernel that handles the physical I/O to a device. This is implemented as a fork process. Same as driver.
I/O interface	<i>See</i> interface.
I/O module	A 16-bit SBB shelf device that integrates the SBB shelf with either an 8-bit single ended, 16-bit single-ended, or 16-bit differential SCSI bus (see SBB).
I/O operation	The process of requesting a transfer of data from a peripheral device to memory (or visa versa), the actual transfer of the data, and the processing and overlaying activity to make both of those happen.
IPI	Intelligent Peripheral Interface. An ANSI standard for controlling peripheral devices by a host computer.
IPI-3 Disk	Intelligent Peripheral Interface Level 3 for Disk
IPI-3 Tape	Intelligent Peripheral Interface Level 3 for Tape
JBOD	Just a bunch of disks. A term used to describe a group of single-device logical units.
kernel	The most privileged processor access mode.
LBN	Logical Block Number.
L_port	A node or fabric port capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.
LED	Light Emitting Diode.
link	A connection between two Fibre Channel ports consisting of a transmit fibre and a receive fibre.
local connection	A connection to the subsystem using either its serial maintenance port or the host SCSI bus. A local connection enables you to connect to one subsystem controller within the physical range of the serial or host SCSI cable.
local terminal	A terminal plugged into the EIA-423 maintenance port located on the front bezel of the controller. <i>See also</i> maintenance terminal.
logical bus	A single-ended bus connected to a differential bus by a SCSI bus signal converter.

logical unit	A physical or virtual device addressable through a target ID number. LUNs use their target bus connection to communicate on the SCSI bus.
logical unit number	LUN. A value that identifies a specific logical unit belonging to a SCSI target ID number. A number associated with a physical device unit during a task I/O operations. Each task in the system must establish its own correspondence between logical unit numbers and physical devices.
logon	Also called login. A procedure whereby a participant, either a person or network connection, is identified as being an authorized network participant.
loop	<i>See</i> arbitrated loop.
loop_ID	A seven-bit value numbered contiguously from zero to 126 decimal and represent the 127 legal AL_PA values on a loop (not all of the 256 hex values are allowed as AL_PA values per FC-AL.)
loop tenancy	The period of time between the following events: when a port wins loop arbitration and when the port returns to a monitoring state.
L_Port	A node or fabric port capable of performing Arbitrated Loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.
LUN	<i>See</i> logical unit number.
LRU	Least recently used. A cache term used to describe the block replacement policy for read cache.
Mbps	Approximately one million (10^6) bits per second—that is, megabits per second.
maintenance terminal	An EIA-423-compatible terminal used with the controller. This terminal is used to identify the controller, enable host paths, enter configuration information, and check the controller status. The maintenance terminal is not required for normal operations. <i>See also</i> local terminal.
member	A container that is a storage element in a RAID array.
metadata	The data written to a disk for the purposes of controller administration. Metadata improves error detection and media defect management for the disk drive. It is also used to support storageset configuration and partitioning. Nontransportable disks also contain metadata to indicate they are uniquely configured for HP StorageWorks environments. Metadata can be thought of as “data about data.”
mirroring	The act of creating an exact copy or image of data.

mirrored write-back caching	A method of caching data that maintains two copies of the cached data. The copy is available if either cache module fails.
mirrorset	<i>See</i> RAID level 1.
MIST	Module Integrity Self-Test.
multibus failover	Allows the host to control the failover process by moving the units from one controller to another.
N_port	A port attached to a node for use with point-to-point topology or fabric topology.
NL_port	A port attached to a node for use in all topologies.
network	In data communication, a configuration in which two or more terminals or devices are connected to enable information transfer.
node	In data communications, the point at which one or more functional units connect transmission lines.
Non-L_Port	A Node of Fabric port that is not capable of performing the Arbitrated Loop functions and protocols. N_Ports and F_Ports loop-capable ports.
non-participating mode	A mode within an L_Port that inhibits the port from participating in loop activities. L_Ports in this mode continue to retransmit received transmission words but are not permitted to arbitrate or originate frames. An L_Port in non-participating mode may or may not have an AL_PA. <i>See also</i> participating mode.
nominal membership	The desired number of mirrorset members when the mirrorset is fully populated with active devices. If a member is removed from a mirrorset, the actual number of members may fall below the “nominal” membership.
node	In data communications, the point at which one or more functional units connect transmission lines. In Fibre Channel, a device that has at least one N_Port or NL_Port.
nonredundant controller configuration	(1) A single controller configuration. (2) A controller configuration that does not include a second controller.
normal member	A mirrorset member that, block-for-block, contains the same data as other normal members within the mirrorset. Read requests from the host are always satisfied by normal members.

normalizing	Normalizing is a state in which, block-for-block, data written by the host to a mirrorset member is consistent with the data on other normal and normalizing members. The normalizing state exists only after a mirrorset is initialized. Therefore, no customer data is on the mirrorset.
normalizing member	A mirrorset member whose contents are the same as all other normal and normalizing members for data that has been written since the mirrorset was created or lost cache data was cleared. A normalizing member is created by a normal member when either all of the normal members fail or all of the normal members are removed from the mirrorset. <i>See also</i> copying member.
NVM	Non-Volatile Memory. A type of memory where the contents survive power loss. Also sometimes referred to as NVMEM.
OCP	Operator control panel. The control or indicator panel associated with a device. The OCP is usually mounted on the device and is accessible to the operator.
offset	A relative address referenced from the base element address. Event Sense Data Response Templates use offsets to identify various information contained within one byte of memory (bits 0 through 7).
other controller	The controller in a dual-redundant pair that is connected to the controller serving the current CLI session. <i>See also</i> this controller.
outbound fiber	One fiber in a link that carries information away from a port.
parallel data transmission	A data communication technique in which more than one code element (for example, bit) of each byte is sent or received simultaneously.
parity	A method of checking if binary numbers or characters are correct by counting the ONE bits. In odd parity, the total number of ONE bits must be odd; in even parity, the total number of ONE bits must be even.
parity bit	A binary digit added to a group of bits that checks to see if errors exist in the transmission.
parity check	A method of detecting errors when data is sent over a communications line. With even parity, the number of ones in a set of binary data should be even. With odd parity, the number of ones should be odd.
parity RAID	<i>See</i> RAIDset.
participating mode	A mode within an L_Port that allows the port to participate in loop activities. A port must have a valid AL_PA to be in participating mode.

partition	A logical division of a container, represented to the host as a logical unit.
PCMCIA	Personal Computer Memory Card Industry Association. An international association formed to promote a common standard for PC card-based peripherals to be plugged into notebook computers. The card commonly known as a PCMCIA card is about the size of a credit card.
PDU	Power distribution unit. The power entry device for HP StorageWorks cabinets. The CDU provides the connections necessary to distribute power to the cabinet shelves and fans.
peripheral device	Any unit, distinct from the CPU and physical memory, that can provide the system with input or accept any output from it. Terminals, printers, tape drives, and disks are peripheral devices.
pluggable	A replacement method that allows the complete system to remain online during device removal or insertion. The system bus must be halted, or quiesced, for a brief period of time during the replacement procedure. See also hot-pluggable.
point-to-point connection	A network configuration in which a connection is established between two, and only two, terminal installations. The connection may include switching facilities.
port	(1) In general terms, a logical channel in a communications system. (2) The hardware and software used to connect a host controller to a communications bus, such as a SCSI bus or serial bus. Regarding the controller, the port is (1) the logical route for data in and out of a controller that can contain one or more channels, all of which contain the same type of data. (2) The hardware and software that connects a controller to a SCSI device.
port_name	A 64-bit unique identifier assigned to each Fibre Channel port. The Port_Name is communicated during the login and port discovery process.
preferred address	The AL_PA which an NL_Port attempts to acquire first during initialization.
primary cabinet	The primary cabinet is the subsystem enclosure that contains the controllers, cache modules, external cache batteries, and the PVA module.

private NL_Port	An NL_Port which does not attempt login with the fabric and only communicates with NL_Ports on the same loop.
program card	The PCMCIA card containing the controller operating software.
protocol	The conventions or rules for the format and timing of messages sent and received.
PTL	Port-Target-LUN. The controller method of locating a device on the controller device bus.
PVA module	Power Verification and Addressing module.
quiesce	The act of rendering bus activity inactive or dormant. For example, “quiesce the SCSI bus operations during a device warm-swap.”
RAID	Redundant Array of Independent Disks. Represents multiple levels of storage access developed to improve performance or availability or both.
RAID level 0	A RAID storageset that stripes data across an array of disk drives. A single logical disk spans multiple physical disks, enabling parallel data processing for increased I/O performance. While the performance characteristics of RAID level 0 is excellent, this RAID level is the only one that does not provide redundancy. Raid level 0 storagesets are sometimes referred to as stripesets.
RAID level 0+1	A RAID storageset that stripes data across an array of disks (RAID level 0) and mirrors the striped data (RAID level 1) to provide high I/O performance and high availability. This RAID level is alternatively called a striped mirrorset.
RAID level 1	A RAID storageset of two or more physical disks that maintain a complete and independent copy of the entire virtual disk's data. This type of storageset has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storagesets are sometimes referred to as mirrorsets.
RAID level 3	A RAID storageset that transfers data parallel across the array disk drives a byte at a time, causing individual blocks of data to be spread over several disks serving as one enormous virtual disk. A separate redundant check disk for the entire array stores parity on a dedicated disk drive within the storageset. <i>See also</i> RAID level 5.
RAID level 5	A RAID storageset that, unlike RAID level 3, stores the parity information across all of the disk drives within the storageset. <i>See also</i> RAID level 3.

RAID level 3/5	A RAID storageset that stripes data and parity across three or more members in a disk array. A RAIDset combines the best characteristics of RAID level 3 and RAID level 5. A RAIDset is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAIDset is sometimes called parity RAID.
RAIDset	<i>See</i> RAID level 3/5.
RAM	Random access memory.
read ahead caching	A caching technique for improving performance of synchronous sequential reads by prefetching data from disk.
read caching	A cache management method used to decrease the subsystem response time to a read request by allowing the controller to satisfy the request from the cache memory rather than from the disk drives.
reconstruction	The process of regenerating the contents of a failed member data. The reconstruct process writes the data to a spareset disk and incorporates the spareset disk into the mirrorset, striped mirrorset, or RAIDset from which the failed member came. <i>See also</i> regeneration.
reduced	Indicates that a mirrorset or RAIDset is missing one member because the member has failed or has been physically removed.
redundancy	The provision of multiple interchangeable components to perform a single function in order to cope with failures and errors. A RAIDset is considered to be redundant when user data is recorded directly to one member and all of the other members include associated parity information.
regeneration	(1) The process of calculating missing data from redundant data. (2) The process of recreating a portion of the data from a failing or failed drive using the data and parity information from the other members within the storageset. The regeneration of an entire RAIDset member is called reconstruction. <i>See also</i> reconstruction.
remote copy	A feature intended for disaster tolerance and replication of data from one storage subsystem or physical site to another subsystem or site. Remote copy also provides methods of performing a backup at either the local or remote site. With remote copy, user applications continue to run while data movement goes on in the background. Data warehousing, continuous computing, and enterprise applications all require remote copy capabilities.

remote copy set	A bound set of two units, one located locally and one located remotely, for long-distance mirroring. The units can be a single disk, or a storageset, mirrorset, or RAIDset. A unit on the local controller is designated as the “initiator” and a corresponding unit on the remote controller is designated as the “target”.
request rate	The rate at which requests are arriving at a servicing entity.
RFI	Radio frequency interference. The disturbance of a signal by an unwanted radio signal or frequency.
replacement policy	The policy specified by a switch with the SET FAILEDSET command indicating whether a failed disk from a mirrorset or RAIDset is to be automatically replaced with a disk from the spareset. The two switch choices are AUTOSPARE and NOAUTOSPARE.
SBB	HP StorageWorks building block. (1) A modular carrier plus the interface required to mount the carrier into a standard HP StorageWorks shelf. (2) any device conforming to shelf mechanical and electrical standards installed in a 3.5-inch or 5.25-inch carrier, whether it is a storage device or power supply.
SCSI	Small computer system interface. (1) An ANSI interface standard defining the physical and electrical parameters of a parallel I/O bus used to connect initiators to devices. (2) a processor-independent standard protocol for system-level interfacing between a computer and intelligent devices including hard drives, floppy disks, CD-ROMs, printers, scanners, and others.
SCSI-A cable	A 50-conductor (25 twisted-pair) cable generally used for single-ended, SCSI-bus connections.
SCSI bus signal converter	Sometimes referred to as an adapter. (1) A device used to interface between the subsystem and a peripheral device unable to be mounted directly into the SBB shelf of the subsystem. (2) a device used to connect a differential SCSI bus to a single-ended SCSI bus. (3) A device used to extend the length of a differential or single-ended SCSI bus. <i>See also</i> DOC (DWZZA-On-a-chip) and I/O module.
SCSI device	(1) A host computer adapter, a peripheral controller, or an intelligent peripheral that can be attached to the SCSI bus. (2) Any physical unit that can communicate on a SCSI bus.
SCSI device ID number	A bit-significant representation of the SCSI address referring to one of the signal lines, numbered 0 through 7 for an 8-bit bus, or 0 through 15 for a 16-bit bus. <i>See also</i> target ID number.

SCSI ID number	The representation of the SCSI address that refers to one of the signal lines numbered 0 through 15.
SCSI-P cable	A 68-conductor (34 twisted-pair) cable generally used for differential bus connections.
SCSI port	(1) Software: The channel controlling communications to and from a specific SCSI bus in the system. (2) Hardware: The name of the logical socket at the back of the system unit to which a SCSI device is connected.
serial transmission	A method transmission in which each bit of information is sent sequentially on a single channel rather than simultaneously as in parallel transmission.
signal converter	<i>See</i> SCSI bus signal converter.
single ended I/O module	A 16-bit I/O module. <i>See also</i> I/O module.
single-ended SCSI bus	An electrical connection where one wire carries the signal and another wire or shield is connected to electrical ground. Each signal logic level is determined by the voltage of a single wire in relation to ground. This is in contrast to a differential connection where the second wire carries an inverted signal.
spareset	A collection of disk drives made ready by the controller to replace failed members of a storage set.
storage array	An integrated set of storage devices.
storage array subsystem	<i>See</i> storage subsystem.
storage set	(1) A group of devices configured with RAID techniques to operate as a single container. (2) Any collection of containers, such as stripesets, mirrorsets, striped mirrorsets, and RAIDsets.
storage set expansion	The dynamic expansion of the storage capacity (size) of a unit. A storage container is created in the form of a concatenation set which is added to the existing storage set defined as a unit.
storage subsystem	The controllers, storage devices, shelves, cables, and power supplies used to form a mass storage subsystem.
storage unit	The general term that refers to storage sets, single-disk units, and all other storage devices that are installed in your subsystem and accessed by the host. A storage unit can be any entity that is capable of storing data, whether it is a physical device or a group of physical devices.

StorageWorks	A family of modular data storage products that allow customers to design and configure their own storage subsystems. Components include power, packaging, cabling, devices, controllers, and software. Customers can integrate devices and array controllers in HP StorageWorks enclosures to form storage subsystems.
	HP StorageWorks systems include integrated SBBs and array controllers to form storage subsystems. System-level enclosures to house the shelves and standard mounting devices for SBBs are also included.
stripe	The data divided into blocks and written across two or more member disks in an array.
striped mirrorset	<i>See RAID level 0+1.</i>
stripeset	<i>See RAID level 0.</i>
stripe size	The stripe capacity as determined by $n-1$ times the chunksize, where n is the number of RAIDset members.
striping	The technique used to divide data into segments, also called chunks. The segments are striped, or distributed, across members of the stripeset. This technique helps to distribute hot spots across the array of physical devices to prevent hot spots and hot disks. Each stripeset member receives an equal share of the I/O request load, improving performance.
surviving controller	The controller in a dual-redundant configuration pair that serves its companion devices when the companion controller fails.
switch	A method that controls the flow of functions and operations in software.
synchronous	Pertaining to a method of data transmission which allows each event to operate in relation to a timing signal. <i>See also asynchronous.</i>
tape	A storage device supporting sequential access to variable sized data records.
target	(1) A SCSI device that performs an operation requested by an initiator. (2) Designates the target identification (ID) number of the device.
target ID number	The address a bus initiator uses to connect with a bus target. Each bus target is assigned a unique target address.
this controller	The controller that is serving your current CLI session through a local or remote terminal. <i>See also other controller.</i>

tape inline exerciser	(TILX) The controller diagnostic software to test the data transfer capabilities of tape drives in a way that simulates a high level of user activity.
topology	An interconnection scheme that allows multiple Fibre Channel ports to communicate with each other. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.
transfer data rate	The speed at which data may be exchanged with the central processor, expressed in thousands of bytes per second (kbytes).
transparent failover	Keeps the storage array available to the hosts by allowing the surviving controller of a dual redundant pair to take over total control of the subsystem and is transparent (invisible) to the hosts.
ULP	Upper Layer Protocol.
ULP process	A function executing within a Fibre Channel node which conforms to the Upper Layer Protocol (ULP) requirements when interacting with other ULP processes.
Ultra SCSI	A Fast-20 SCSI bus. <i>See also</i> Wide Ultra SCSI.
unit	A container made accessible to a host. A unit may be created from a single disk drive or tape drive. A unit may also be created from a more complex container such as a RAIDset. The controller supports a maximum of eight units on each target. <i>See also</i> target and target ID number.
unwritten cached data	Sometimes called unflushed data. <i>See</i> dirty data.
UPS	Uninterruptible power supply. A battery-powered power supply guaranteed to provide power to an electrical device in the event of an unexpected interruption to the primary power supply. Uninterruptible power supplies are usually rated by the amount of voltage supplied and the length of time the voltage is supplied.
VHDCI	Very high-density-cable interface. A 68-pin interface. Required for Ultra-SCSI connections.
virtual terminal	A software path from an operator terminal on the host to the controller's CLI interface, sometimes called a host console. The path can be established via the host port on the controller or via the maintenance port through an intermediary host.
VTDPY	An abbreviation for Virtual Terminal Display Utility.

warm swap	A device replacement method that allows the complete system to remain online during device removal or insertion. The system bus may be halted, or quiesced, for a brief period of time during the warm-swap procedure.
Wide Ultra SCSI	Fast/20 on a Wide SCSI bus.
Worldwide name	A unique 64-bit number assigned to a subsystem by the Institute of Electrical and Electronics Engineers (IEEE) and set by manufacturing prior to shipping. This name is referred to as the node ID within the CLI.
write-back caching	A cache management method used to decrease the subsystem response time to write requests by allowing the controller to declare the write operation “complete” as soon as the data reaches its cache memory. The controller performs the slower operation of writing the data to the disk drives at a later time.
write-through caching	Write-through caching always writes directly to disk, ensuring that the application is never tricked into believing that the data is on the disk when it may not be. This results in highest data integrity, though with slightly reduced performance.
write hole	The period of time in a RAID level 1 or RAID level 5 write operation when an opportunity emerges for undetectable RAIDset data corruption. Write holes occur under conditions such as power outages, where the writing of multiple members can be abruptly interrupted. A battery backed-up cache design eliminates the write hole because data is preserved in cache and unsuccessful write operations can be retried.
write-through cache	A cache management technique for retaining host write requests in read cache. When the host requests a write operation, the controller writes data directly to the storage device. This technique allows the controller to complete some read requests from the cache, greatly improving the response time to retrieve data. The operation is complete only after the data to be written is received by the target storage device. This cache management method may update, invalidate, or delete data from the cache memory accordingly, to ensure that the cache contains the most current data.

index

A

ADD CONNECTIONS

- multiple-bus failover [48](#)
- transparent failover [46](#)

ADD UNIT

- multiple-bus failover [48](#)
- transparent failover [46](#)

adding

- virtual disks [228](#)

adding a disk drive to the spareset

- configuration options [160](#)

adding disk drives

- configuration options [160](#)

Agent

- functions [124](#)

array of disk drives

- [77](#)

assigning unit numbers

assignment

- unit numbers

- fabric topology [159](#)

- unit qualifiers

- fabric topology [159](#)

assignment of unit numbers

- fabric topology

- partition [159](#)

- single disk [159](#)

asynchronous event service

- [228](#)

audience

- [12](#)

authorized reseller, HP

- [21](#)

autospare

- enabling

- fabric topology [161](#)

availability

84

B

Back up, Clone, Move Data

- [177](#)

backup

- cloning data [179](#)

- subsystem configuration [178](#)

C

cabling

- controller pair [146](#)

- multiple-bus failover

- fabric topology configuration [145](#)

- single controller [136](#)

cache modules

- location [28, 29](#)

- read caching [37](#)

- write-back caching [37](#)

- write-through caching [38](#)

caching techniques

- mirrored [39](#)

- read caching [37](#)

- read-ahead caching [37](#)

- write-back caching [37](#)

- write-through caching [38](#)

changing switches

- configuration options [163](#)

chunk size

- choosing for RAIDsets and stripesets [93](#)

- controlling stripesize [93](#)

- using to increase request rate [94](#)

- using to increase write performance [95](#)

CHUNKSIZE 93
CLI configuration
 example 168
CLI configurations 165
CLI prompt
 changing
 fabric topology 160
Client
 removing 226
 uninstalling 226
CLONE and SNAPSHOT 183
CLONE utility
 backup 179
cloning
 backup 179
command console LUN 39
 SCSI-2 mode 49
 SCSI-3 mode 49
comparison of container types 77
configuration
 backup 178
 fabric topology
 devices 152, 153
 multiple-bus failover
 cabling 145
 multiple-bus failover using CLI 168
 single controller cabling 135
 restoring 95
 rules 65
configuration options
 fabric topology
 adding a disk drive to the spareset 160
 adding disk drives 160
 changing switches
 device 163
 displaying the current switches 163
 initialize 163
 RAIDset and mirrorset 163
 unit 163
 changing switches for a storageset or
 device 162
 changing the CLI prompt 160
 deleting a storageset 162
 enabling autospare 161
 removing a disk drive from the spareset
 161
 configuring
 pager notification 228
 configuring devices
 fabric topology 152, 153
 connections 41
 naming 41
containers
 attributes 76
 comparison 77
 illustrated 76, 153
 mirrorsets 82
 planning storage 76
 stripesets 81
controllers
 cabling 136, 146
 location 28, 29
 node IDs 58
 worldwide names 58
conventions
 document 25
 equipment symbols 17
 text symbols 25
creating
 storageset and device profiles 78
Creating Clones for Backup 179

D

Destroy/Nodestroy parameters 96
device switches
 changing
 fabric topology 163
devices
 changing switches
 fabric topology 162
configuration
 fabric topology 152, 153
 creating a profile 78
disk drives

adding
 fabric topology 160
adding to the spareset
 fabric topology 160
array 77
corresponding storage sets 98
dividing 88
removing from the spareset
 fabric topology 161
displaying the current switches
 fabric topology 163
dividing storage sets 88
document
 conventions 25
 prerequisites 12
 related documentation 12

E

enabling
 switches 90
equipment symbols 17
erasing metadata 96
establishing a local connection 134

F

fabric topology
 configuration
 single controller cabling 135
failover 31
 multiple-bus 34
 transparent 31
First enclosure of multiple-enclosure subsystem
 storage map template 1 198, 201, 203,
 205, 208, 210, 213

functions
 Agent 124

G

geometry
 initialize switches 96
Geometry parameters 96

getting help 20

H

HACMP 229
help, obtaining 25
Host access
 restricting in multiple-bus failover mode
 disabling access paths 55
host access
 restricting by offsets
 multiple-bus failover 57
 transparent failover 55
 restricting in multiple-bus failover mode 55
 restricting in transparent failover mode 51
 disabling access paths 53
 separate links 52

host adapter

 installation 108
 preparation 108

host connections 41

 naming 41

Host storage configuration
 verify 164

HP

 authorized reseller 21
 storage web site 20
 technical support 20

HSG Agent
 install and configure 123
 network connection 127
 overview 126

I

initialize switches
 changing
 fabric topology 163
 CHUNKSIZE 93
 geometry 96
 NOSAVE_CONFIGURATION 95
 SAVE_CONFIGURATION 95
Insight Manager 228

installation
 invalid network port assignments 222
 there is no disk in the drive message 223
integrating SWCC 228
invalid network port assignments 222

J

JBOD 77

L

LOCATE
 find devices 98

location
 cache module 28, 29
 controller 28, 29

LUN IDs
 general description 60

LUN presentation 47

M

maintenance port connection
 establishing a local connection 134
 illustrated 134

mapping storagesets 98

messages
 there is no disk in the drive 223

mirrored caching
 enabling 39
 illustrated 39

mirrorset switches
 changing
 fabric topology 163

mirrorsets
 planning considerations 82
 important points 83
 switches 91

Model 2200 Storage Maps
 examples 69

moving storagesets 191

multiple-bus failover
 restricting host access by offsets 57

multiple-bus failover 34
ADD CONNECTIONS command 48
ADD UNIT command 48
ADD UNITcommand 48
CLI configuration procedure
 fabric topology 168
fabric topology configuration
 cabling 145
host connections 48
restricting host access 55
 disabling access paths 55
SET CONNECTIONS command 48
SET UNITcommand 48

N

network port assignments 222
node IDs 58
restoring 59
NODE_ID
 worldwide name 58
NOSAVE_CONFIGURATION 95

O

offset
 LUN presentation 47
restricting host access
 multiple-bus failover 57
 transparent failover 55
SCSI version factor 48
online help
 SWCC 228
options
 for mirrorsets 91
 for RAIDsets 91
 initialize 93
other controller 29

P

pager notification 228
 configuring 228
partitions

- assigning a unit number
 - fabric topology 159
- defining 89
- planning considerations 88
 - guidelines 89
- performance 84
- Physical connection, making 108
- planning 61
 - overview 78
 - striped mirrorsets 87
 - stripesets 81
- Planning a subsystem 27
- planning configurations
 - where to start 62
- planning considerations 84
- planning storage
 - containers 76
- planning storagesets
 - characteristics
 - changing switches 90
 - enabling switches 90
 - initialization switch 90
 - storageset switch 90
 - unit switch 90
 - switches
 - initialization 93
 - storageset 91
- prerequisites 12
- profiles
 - creating 78
 - description 78
 - storageset 195
 - example 196
- R**
 - rack stability, warning 19
- RAIDset switches
 - changing
 - fabric topology 163
- RAIDsets
 - choosing chunk size 93
 - maximum membership 86
- planning considerations 84
 - important points 85
 - switches 91
- read caching
 - enabled for all storage units 37
 - general description 37
- read requests
 - decreasing the subsystem response time with
 - read caching 37
- read-ahead caching 37
 - enabled for all disk units 37
- related documentation 12
- removing
 - Client 226
- request rate 94
- requirements
 - host adapter installation 108
- restricting host access
 - disabling access paths
 - multiple-bus failover 55
 - transparent failover 53
 - multiple-bus failover 55
 - separate links
 - transparent failover 52
 - transparent failover 51
- S**
 - SAVE_CONFIGURATION 95
 - saving configuration 95
 - SCSI version
 - offset 48
 - SCSI-2
 - assigning unit numbers 49
 - command console lun 49
 - SCSI-3
 - assigning unit numbers 49
 - command console lun 49
 - Second enclosure of multiple-enclosure subsystem
 - storage map template 2 199
 - selective storage presentation 51
 - SET CONNECTIONS

- multiple-bus failover [48](#)
- transparent failover [46](#)
- SET UNIT
 - multiple-bus failover [48](#)
- setting
 - controller configuration handling [95](#)
- single disk (JBOD)
 - assigning a unit number
 - fabric topology [159](#)
- Single-enclosure subsystem
 - storage map template [1 198](#)
- software installation
 - new [110](#)
 - upgrade [110](#)
- storage
 - creating map [98](#)
 - profile
 - example [196](#)
- storage configurations [61](#)
- storage map [98](#)
- Storage map template [1 198](#)
 - first enclosure of multiple-enclosure subsystem [198](#)
 - single enclosure subsystem [198](#)
- Storage map template [2 199](#)
 - second enclosure of multiple-enclosure subsystem [199](#)
- Storage map template [3 200](#)
 - third enclosure of multiple-enclosure subsystem [200](#)
- Storage map template [4](#)
 - first enclosure of multiple-enclosure subsystem [201](#)
- Storage map template [5](#)
 - first enclosure of multiple-enclosure subsystem [203](#)
- Storage map template [6](#)
 - first enclosure of multiple-enclosure subsystem [205](#)
- Storage map template [7](#)
- first enclosure of multiple-enclosure subsystem [208](#)
- Storage map template [8](#)
 - first enclosure of multiple-enclosure subsystem [210](#)
- Storage map template [9](#)
 - first enclosure of multiple-enclosure subsystem [213](#)
- storageset
 - deleting
 - fabric topology [162](#)
 - fabric topology
 - changing switches [162](#)
 - planning considerations [80](#)
 - mirrorsets [82](#)
 - partitions [88](#)
 - RAIDsets [84](#)
 - striped mirrorsets [86](#)
 - stripesets [80](#)
 - profile [78](#)
 - profiles [195](#)
- storageset profile [78](#)
- storageset switches
 - SET command [91](#)
- storagesets
 - creating a profile [78](#)
 - moving [191](#)
- striped mirrorsets
 - planning [87](#)
 - planning considerations [86](#)
- stripesets
 - distributing members across buses [82](#)
 - planning [81](#)
 - planning considerations [80](#)
 - important points [81](#)
- subsystem
 - saving configuration [95](#)
- subsystem configuration
 - backup [178](#)
- SWCC [124](#)
 - additional information [228](#)
 - integrating [228](#)

- online help 228
 - switches
 - changing 90
 - changing characteristics 90
 - CHUNKSIZE 93
 - enabling 90
 - mirrorsets 91
 - NOSAVE_CONFIGURATION 95
 - RAIDset 91
 - SAVE_CONFIGURATION 95
 - switches for storagesets
 - overview 90
 - symbols in text 25
 - symbols on equipment 17
- T**
- technical support, HP 20
 - templates
 - subsystem profile 195
 - terminology
 - other controller 29
 - this controller 29
 - text symbols 25
- Third enclosure of multiple-enclosure subsystem
 - storage map template 3 200
 - this controller 29
 - transparent failover 31
 - ADD CONNECTIONS 46
 - ADD UNIT 46
 - matching units to host connections 46
 - restricting host access 51
 - disabling access paths 53
 - separate links 52
 - restricting host access by offsets 55
 - SET CONNECTIONS 46
 - troubleshooting
 - invalid network port assignments 222
 - there is no disk in the drive message 223
- U**
- uninstalling
- Client 226
 - unit numbers
 - assigning 46
 - fabric topology 159
 - assigning depending on SCSI version 49
 - assigning in fabric topology
 - partition 159
 - single disk 159
 - unit qualifiers
 - assigning
 - fabric topology 159
 - unit switches
 - changing
 - fabric topology 163
 - units
 - LUN IDs 60
 - Upgrade procedures
 - solution software 121
- V**
- virtual disks
 - adding 228
- W**
- warning
 - rack stability 19
 - symbols on equipment 17
 - web sites
 - HP storage 20
 - where to start 27
 - worldwide names 58
 - NODE_ID 58
 - REPORTED PORT_ID 58
 - restoring 59
 - write performance 95
 - write requests
 - improving the subsystem response time with
 - write-back caching 37
 - placing data with write-through caching 38
 - write-back caching
 - general description 37

write-through caching

general description 38